

10/19/00

Please type a plus sign (+) inside this box

+

10/20/00

PTO/SB/05 (4/98)

UTILITY PATENT APPLICATION TRANSMITTAL

(Only for nonprovisional applications under 37 CFR § 1.53(b))

Attorney Docket No.

294438005US1

First Inventor or Application Identifier

Ken Abbott

Title

INTERFACE FOR EXCHANGING CONTEXT DATA

Express Mail Label No.

EL696997286US

APPLICATION ELEMENTS

See MPEP chapter 600 concerning utility patent application contents.

ADDRESS TO:

Box Patent Application
Assistant Commissioner for Patents
Washington, D.C. 202311. ☐ Authorization for Extensions & Fee Transmittal
(Submit an original and a duplicate for fee processing)5. ☐ Microfiche Computer Program (Appendix)2. ☒ Specification [Total Pages] **66**
(preferred arrangement set forth below)6. Nucleotide and/or Amino Acid Sequence Submission
(if applicable, all necessary)

- Descriptive Title of the Invention
- Cross References to Related Applications
- Statement Regarding Fed sponsored R & D
- Reference to Microfiche Appendix
- Background of the Invention
- Brief Summary of the Invention
- Brief Description of the Drawings (if filed)
- Detailed Description
- Claim(s)
- Abstract of the Disclosure

a. ☐ Computer-Readable Copyb. ☐ Paper Copy (identical to computer copy)c. ☐ Statement verifying identity of above copy3. ☒ Drawing(s) (35 USC 113) [Total Sheets] **10**4. Oath or Declaration [Total Pages] **1**a. ☐ Newly executed (original or copy)b. ☐ Copy from a prior application (37 CFR 1.63(d))
(for continuation/divisional with Box 16 completed)

- i. ☐ **DELETION OF INVENTOR(S)**
Signed statement attached deleting
inventor(s) named in the prior application,
see 37 CFR 1.63(d)(2) and 1.33(b)

*NOTE FOR ITEMS 1 & 13: IN ORDER TO BE ENTITLED TO
PAY SMALL ENTITY FEES, A SMALL ENTITY STATEMENT IS
REQUIRED (37 C.F.R. § 1.27), EXCEPT IF ONE FILED IN A
PRIOR APPLICATION IS RELIED UPON (37 C.F.R. § 1.28).

ACCOMPANYING APPLICATION PARTS

7. ☐ Assignment Papers (cover sheet & document(s))8. ☐ 37 CFR 3.73(b) Statement ☐ Power of Attorney
(when there is an assignee)9. ☐ English Translation Document (if applicable)10. ☐ Information Disclosure Statement (IDS)/PTO-1449 ☐ Copies of IDS Citations11. ☐ Preliminary Amendment12. ☒ Return Receipt Postcard12. ☐ Small Entity ☐ Statement filed in prior application,
Statement(s) Status still proper and desired14. ☐ Certified Copy of Priority Document(s)
(if foreign priority is claimed)15. ☐ Other: _____

16. If a CONTINUING APPLICATION, check appropriate box and supply the requisite information below and in a preliminary amendment

☐ Continuation ☐ Divisional ☒ Continuation-In-Part (CIP) of prior Application No.: **09/541,328, 09/541,326 and 09/216,193**

Prior application information: Examiner _____ Group / Art Unit _____

For CONTINUATION or DIVISIONAL apps only: The entire disclosure of the prior application, from which an oath or declaration is supplied under Box 4b, is considered a part of the disclosure of the accompanying continuation or divisional application and is hereby incorporated by reference. The incorporation can only be relied upon when a portion has been inadvertently omitted from the submitted application parts.

☐ Claims the benefit of Provisional Application No. _____

17. CORRESPONDENCE ADDRESS

Customer Number 25096 / Barcode



25096

PATENT TRADEMARK OFFICE

Respectfully submitted,

TYPED or PRINTED NAME Steven D. LawrenzREGISTRATION NO. 37,376SIGNATURE [Signature]Date 10/19/00

INTERFACE FOR EXCHANGING CONTEXT DATA

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a continuation-in-part of U.S. Patent Application No. 09/541,328 entitled "INTERFACE FOR EXCHANGING CONTEXT DATA" filed April 2, 2000, U.S. Patent Application No. 09/541,326, entitled "LOGGING AND ANALYZING COMPUTER USER'S CONTEXT DATA" filed April 2, 2000 and U.S. Patent Application No. 09/216,193, entitled "METHOD AND SYSTEM FOR CONTROLLING PRESENTATION OF INFORMATION TO A USER BASED ON THE USER'S CONDITION" filed December 18, 1998, which are hereby incorporated by reference in their entirety.

TECHNICAL FIELD

The present invention is directed to the field of context modeling and, more particularly, to the field of data exchange for context modeling.

BACKGROUND

Wearable computers are devices that commonly serve as electronic companions and intelligent assistants to their users. A wearable computer is typically strapped to its user's body or carried by its user in a holster, and can contain a variety of both input and output devices. A wearable computer can output information to its user using, for example, display eyeglasses, audio speakers, or a tactile output device. A wearable computer can receive instructions and other input from its user via input devices such as a keyboard, various pointing devices, or an audio microphone. A wearable computer can receive information about its surroundings using sensors, such as barometric pressure and temperature sensors, global positioning system devices, or a heart rate monitor for determining the heart rate of its user and can receive additional information via communication devices, such as various types of network connections. A

wearable computer can exchange information with other devices using communication schemes such as infrared communication, radio communication, or cellular modems.

Many applications for wearable computers utilize data received by the wearable computer from sensors. For example, a position mapping application for a
5 wearable computer may utilize data received from a global positioning system device to plot its user's physical location and determine whether that position is within a special region. In this example, the global positioning system device produces data that is consumed by the position mapping application.

In conventional wearable computer systems, the position mapping
10 application would be designed to interact directly with the global positioning system device to obtain the needed data. For example, the application may be required to instruct the device to obtain position information, retrieve the information obtained by the device, convert it to conventional latitude and longitude representation, and determine whether the represented location is within the special region.

Such direct interaction between applications and sensors to obtain and
15 process data has significant disadvantages. First, developing an application to interact directly with a particular sensor can introduce into the application dependencies on that sensor. Accordingly, the application may need to be subsequently modified to interact successfully with alternatives to that sensor provided by other manufacturers, or even to interact successfully with future versions of the sensor from the same manufacturer.
20

Second, direct interaction between the application and the sensor can give
rise to conflicts between multiple applications that consume the same data. For example, if the position mapping application was executing on the same wearable computer as a
25 second application for determining the user's distance from home that also used the global positioning system device, the two applications' interactions with the device could interfere with one another.

Third, direct interaction between the application and the sensor can give
rise to conflicts between multiple sensors that produce the same data. For example, if the position mapping application was executing on a wearable computer that had access to
30 both the global positioning system device and an indoor positioning system, the application might well have trouble determining which device to use to determine the user's current position, and/or have trouble reconciling data produced by both devices.

Fourth, performing the derivation of abstract information from data observable by sensors in each application that requires the derived abstract information necessitates redundant functionality in each such application, and precludes the development of multiple competing algorithms to derive the same abstract information. rather than having to themselves process data from the sensor to derive more abstract information from data observable by sensors, it would be advantageous for applications to be able to rely on a separate programmatic entity to derive such abstract information. For example, it would be more convenient for the position mapping application to be able to rely on a separate programmatic entity to determine whether the user is in a special region based upon the user's location. It would further be advantageous for such applications to share a single separate programmatic entity, rather each implementing the same derivation functionality.

Accordingly, a facility for exchanging information between sensors and applications in a wearable computer system would have significant utility.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 illustrates an embodiment of the characterization module which executes on a general-purpose body-mounted wearable computer 120 worn by a user 110.

Figure 2 illustrates an exemplary computer system 200 on which an embodiment of the characterization module is executing.

Figure 3 is a data flow diagram showing a sample exchange of attributes performed by the characterization module.

Figure 4 is a data structure diagram showing a context server table that maintains a portion of the state of the characterization module.

Figure 5 is a data structure diagram showing an attribute instance table that maintains a portion of the state of the characterization module.

Figure 6 is a data structure diagram showing a context client table that maintains a portion of the state of the characterization module.

Figure 7 is a data structure diagram showing an attribute or instance registration table that maintains a portion of the state of the characterization module.

Figure 8 is a flow diagram showing the steps preferably performed by the characterization module when the GetAttribute function is called.

Figure 9 is a data structure diagram showing updated contents of the attribute instance table.

Figure 10 is a data structure diagram showing a condition table that contains a portion of the state of the characterization module.

Figure 11 is a data structure diagram showing a condition monitor table that maintains a portion of the state of the characterization module.

Figure 12 is a data structure diagram showing updated contents of the condition monitor table.

Figure 13 is a data flow diagram showing the operation of the facility without a characterization module.

Figure 14 is a data structure diagram showing an attribute instance property table that maintains a portion of the state of the characterization module.

DETAILED DESCRIPTION

A software facility for exchanging information between sources of context data and consumers of context data (“the facility”) is provided. In a preferred embodiment, a characterization module operating in a wearable computer system receives context information, in the form of individual attributes each modeling an aspect of the wearable computer system, its user, or the surrounding environment, from one or more context servers, and provides it to one or more context clients. The facility reduces dependencies of context client applications on specific sensors and other input devices, reduces conflicts between context client applications that consume the same context data, resolves conflicts between multiple sensors or other input devices that produce the same data, isolates the derivation of derived attributes from context client applications, and therefore obviates the redundant effort of implementing the derivation of derived attributes in each context server, and facilitates the development and use of competing algorithms to derive derived attributes.

Attributes represent measures of specific context elements such as ambient temperature, location and current user task. Each attribute preferably has the following

properties: a name, a value, an uncertainty level, units, and a time stamp. Attributes provided through the characterization module by a context server may either be “measured,” in that they are directly received from an input device, or “derived,” in that they are the result of performing processing on values directly obtained from input devices other attributes. Indeed, a derived attribute may be produced by performing additional processing on other derived attributes. Context servers, in addition to providing attributes through the characterization module, may also provide other functions. For example, an application, such as an electronic mail application, could serve as a context server by providing attributes through the characterization module. In addition to the source of attributes described above, such an “expanded” context server may provide attributes relating to the other functions of the expanded context server. For example, an electronic mail application context server could provide an attribute indicating other new messages are waiting. Indeed, the same program module may operate both as a context client and a context server.

Two or more different context servers may preferably supply to the characterization module their own values for a single attribute. For example, a first context server can supply a value for a user.location attribute based on data received from a global positioning system device, while a second context server can supply a value for the user.location attribute based on data received from an indoor positioning device. These separate values for the same attribute are referred to as separate “instances” of the attribute. The characterization module preferably provides a variety of different approaches, called “mediators,” for determining, when an attribute requested by a context client has more than one instance, what attribute value to provide in response to the attribute request.

When the characterization module obtains an attribute value from a context server, it preferably caches it for responding to future requests for the attribute from context clients. Such attribute requests may specify a specific instance of the attribute—that is, a specific context server from which the attribute is to be obtained—or may specify that the attribute is to be obtained by applying a particular mediator to whatever instances of the attribute are available, or may utilize a default mediator to mediate between available instances of the attribute. When the characterization module receives an attribute request from a context client, it identifies the attribute instances available to

satisfy the request, and, for each, determines whether the value for the attribute instance cached by the characterization module was obtained sufficiently recently from the corresponding context server. If not, the characterization module requests an updated value for the attribute instance from the corresponding context server before performing any necessary mediation and returning a value to the context client.

Context servers and context clients preferably interact with the characterization module by calling functions provided by the characterization module. These functions are said to collectively comprise an “application programming interface” (or “API”) to the characterization module.

A context client that seeks to request a particular attribute or attribute instance may call a LaunchContextServer function to launch a particular context server that is not executing. Each executing context server preferably registers with the characterization module by calling a RegisterContextServer function, identifying itself. The context server goes on to make an instance of an attribute available to context clients via the characterization module by following a CreateAttributeInstance function. A particular context server may preferably provide a number of different attributes by calling the CreateAttributeInstance function multiple times. Before seeking to consume an attribute, a context client calls a RegisterContextClient function, identifying itself. In order to consume an attribute, a context client calls a RegisterToConsumeAttributeOrInstance function, identifying itself and an attribute that it seeks to consume. To help determine which attributes to consume, a context client may call a EnumerateAttributes function to obtain a list of the attributes available from the characterization module. In order to actually retrieve an attribute value, a context client calls a GetAttribute function, identifying the attribute and any attribute processing that should be applied, such as a specific type of mediation between different instances of the attribute obtained from different context servers. For attributes that have multiple instances in the characterization module, a context client may call a GetAllAttributeInstances function to obtain a value of the attribute for each attribute instance. To force a particular context server to reevaluate all of its attribute instances, a context client may call a CompleteContextServerEvaluation function. Also, to retrieve attributes reflecting aspects of the configuration of the characterization module, a context client or other program may call a GetCharacterizationModuleAttribute function. A

context client that consumes a particular attribute may create a condition in the characterization module that tests that attribute by calling a CreateCondition function. Once a context client has created a condition, it can evaluate the condition by calling an EvaluateCondition function, identifying the condition. Once a context client has created
5 a condition, it may go on to call a CreateConditionMonitor function to create a condition monitor that monitors the condition and notifies the context server when the condition is satisfied. To suspend operation of a created condition monitor, a context server may call a StopConditionMonitor function, and to resume its operation, may call a StartConditionMonitor function. The context server may remove a condition monitor that
10 it created by calling a RemoveConditionMonitor function, and, correspondingly, may remove a condition that it created by calling a RemoveCondition function. A context client may unregister to consume a particular attribute by calling an UnregisterToConsumeAttributeOrInstance function. A context client may unregister itself with the characterization module by calling an UnregisterContextClient function. A
15 context server may – though need not – remove attribute instances that it has registered by calling a RemoveAttributeInstance function. Before it does, however, it may – though need not – first call a CheckAttributeInstanceDependencies function to determine whether any context clients currently depend upon that attribute instance. Once it has removed its attribute instances, a context server may unregister with the characterization
20 module by calling an UnregisterContextServer function. These API functions are discussed in greater detail below in conjunction with an example.

Figure 1 illustrates an embodiment of the characterization module which executes on a general-purpose body-mounted wearable computer 120 worn by a user 110. Many wearable computers are designed to act as constant companions and intelligent
25 assistants to a user, and are often strapped to a user's body or mounted in a holster. The computer system may also be incorporated in the user's clothing, be implanted in the user, follow the user, or otherwise remain in the user's presence. In one preferred embodiment the user is human, but in additional preferred embodiments, the user may be an animal, a robot, a car, a bus, or another entity whose context data is to be processed.
30 Indeed, the computer system may have no identifiable user, but rather operate as an independent probe, processing context data in an arbitrary location. The wearable computer 120 has a variety of user-worn user input devices including a microphone 124,

a hand-held flat panel display 130 with character recognition capabilities, and various other user input devices 122. Similarly, the computer has a variety of user-worn output devices that include the hand-held flat panel display, an earpiece speaker 132, an eyeglass-mounted display 134, and a tactile display 136. In addition to the various user-worn user input devices, the computer can also receive information from various user sensor input devices 116 and from environment sensor input devices 128, including video camera 121. The characterization module can receive and process the various input information received by the computer, such as from context servers that process the input information and generate attribute values, and can present information to the user on the various output devices accessible to the computer.

In the current environment, computer 120 is accessible to a computer 150 (e.g., by being in line-of-sight wireless proximity or by being reachable via a long-distance communication device such as a cellular phone) which also has a variety of input and output devices. In the illustrated embodiment the computer 150 is non-portable, although the body-mounted computer of the user can similarly communicate with a variety of other types of computers, including body-mounted computers of other users. The devices from which the non-portable computer can directly receive information include various user input devices 152 and various user sensor input devices 156. The non-portable computer can output information directly to a display 160, a speaker 162, an olfactory device 164, and a printer 166. In the illustrated embodiment, the body-mounted computer can communicate with the non-portable computer via a wireless transmission medium. In this manner, the characterization module can receive information from the user input devices 152 and the user sensor devices 156 after the information has been transmitted to the non-portable computer and then to the body-mounted computer. Alternately, the body-mounted computer may be able to directly communicate with the user input devices 152 and the user sensor devices 156, as well as with other various remote environment sensor input devices 158, without the intervention of the non-portable computer 150. Similarly, the body-mounted computer may be able to supply output information to the display 160, the speaker 162, the olfactory device 164, and the printer 166, either directly or via the non-portable computer, and directly to the telephone 168. As the user moves out of range of the remote input and output devices, the

characterization module will be updated to reflect that the remote devices are not currently available.

The various input devices allow the characterization module or an application such as a context server (not shown) executing on the computer system 120 to monitor the user and the environment and to model their current condition. Such a model can be used by various applications, such as context clients, for various purposes. A model of the current conditions can include a variety of condition variables that represent information about the user, the computer, and the user's environment at varying levels of abstraction. For example, information about the user at a low level of abstraction can include raw physiological data (*e.g.*, heart rate and EKG) and geographic information (*e.g.*, location and speed), while higher levels of abstraction may attempt to characterize or predict the user's physical activity (*e.g.*, jogging or talking on a phone), emotional state (*e.g.*, angry or puzzled), desired output behavior for different types of information (*e.g.*, to present private family information so that it is perceivable only to myself and my family members), and cognitive load (*i.e.*, the amount of attention required for the user's current activities). Background information which changes rarely or not at all can also be included, such as the user's age, gender and visual acuity. The model can similarly hold environment information at a low level of abstraction, such as air temperature or raw data from a motion sensor, or at higher levels of abstraction, such as the number and identities of nearby people, objects, and locations. The model of the current conditions can additionally include information added explicitly from other sources (*e.g.*, application programs), as well as user-specified or system-learned defaults and preference information.

Those skilled in the art will appreciate that computer systems 120 and 150, as well as their various input and output devices, are merely illustrative and are not intended to limit the scope of the present invention. The computer systems may contain additional components or may lack some illustrated components. For example, it is possible that the characterization module can be implemented on the non-portable computer, with the body-mounted computer replaced by a thin context client such as a transmitter/receiver for relaying information between the body-mounted input and output devices and the non-portable computer. Alternately, the user may not wear any devices or computers.

In addition, the body-mounted computer may be connected to one or more networks of other devices through wired or wireless communication means (*e.g.*, wireless RF, a cellular phone or modem, infrared, physical cable, a docking station, etc.), either with or without support from other computers such as the computer 150. For example, the body-mounted computer of a user can make use of output devices in a smart room, such as a television and stereo when the user is at home, if the body-mounted computer can transmit information to those devices via a wireless medium or if a cabled or docking mechanism is available. Alternately, kiosks or other information devices can be installed at various locations (*e.g.*, in airports or at tourist spots) to transmit relevant information to body-mounted computers within the range of the information device.

Those skilled in the art will also appreciate that specialized versions of the body-mounted computer and characterization module can be created for a variety of purposes.

Figure 2 illustrates an exemplary computer system 200 on which an embodiment of the characterization module is executing. The computer includes a memory 230, a CPU 210, a persistent storage device 250 such as a hard drive, and input/output devices including a microphone 222, a video camera 223, a computer-readable media drive 224, such as a CD-ROM drive, a visual display 225, a speaker 226, and other devices 228. The memory preferably includes the characterization module 231, as well as information reflecting the current state of the characterization module (characterization module state) 232. The memory further contains software modules 233, 234, and 237 that consume attributes and are therefore context clients, and software modules 235, 236, and 237 which provide attributes and are therefore context servers. While items 231-237 are preferably stored in memory while being used, those skilled in the art will appreciate that these items, or portions of them, can be transferred between memory and the persistent storage device for purposes of memory management and data integrity.

The facility preferably utilizes a plain-language, hierarchical, taxonomic attribute nomenclature to name attributes. The attribute names within the nomenclature are preferably specific so that there is no ambiguity as to what they represent. The facility preferably supports the extension of the nomenclature by adding new attribute names that conform to the hierarchical taxonomy of the nomenclature. The nomenclature

preferably has attribute names relating to the user, such as user.position, user.speed, and user.direction, providing information about the user's position, speed, and direction, respectively. The nomenclature preferably contains attribute names for various user moods, such as user.mood.happiness, user.mood.anger, and user.mood.confusion. The nomenclature preferably includes attribute names for user activities, such as user.activity.driving, user.activity.eating, and user.activity.sleeping. The nomenclature preferably includes attribute names for user physiology values, such as user.physiology.pulse, user.physiology.body_temperature, and user.physiology.blood_pressure. The nomenclature preferably includes attribute names for similar attributes of people other than the user, such as person.John_Smith.mood.happiness. The nomenclature preferably includes attribute names for aspects of the computer system or "platform." For example, the nomenclature preferably includes attribute names for aspects of the platform's user interface, such as platform.user_interface.oral_input_device_availability and platform.user_interface.visual_output_device_availability. The nomenclature preferably includes attribute names for attributes relating to the central processing unit, or "CPU," of the platform, such as platform.cpu_load and platform.clock_speed. The nomenclature preferably also provides attribute names for aspects of the local environment, such as environment.local.time, environment.local.temperature, and environment.local.ambient_noise_level. The nomenclature preferably also includes attribute names for remote environments, such as environment.place.chicago.time and environment.place.san_diego.temperature. The nomenclature preferably further provides attribute names relating to specific applications. For example, the nomenclature preferably provides attribute names for aspects of an electronic mail application, such as application.mail.available, application.mail.new_messages_waiting, and application.mail.messages_waiting_to_be_sent. In this manner, the attribute nomenclature used by the facility provides effective names for attributes relating to the user, the computer system, and the environment. Additional detail on the attribute nomenclature utilized by the facility is provided in U.S. Patent Application No. _____, entitled "Supplying User Context Data With Dynamically Specified Suppliers and Consumers," which is hereby incorporated by reference in its entirety.

Figure 3 is a data flow diagram showing a sample exchange of attributes performed by the characterization module. The diagram shows characterization module 300, as well as five other software modules, 310, 320, 330, 340, and 350. Software modules 310, 320, and 330 are said to be context servers, in that they provide attributes to the characterization module. For example, context server 330 provides an user.in_region attribute 331. It can be seen that context servers may provide more than one attribute. For example, context server 320 provides a user.location attribute 321 and an user.elevation attribute 322. It can further be seen that a single attribute may be provided by more than one context server. For example, context server 310 provides user.location attribute 311, while context server 320 provides user.location attribute 321. The characterization module preferably provides functionality for mediating between these two separate instances of a user.location attribute when the user.location attribute is requested by a context client.

Software modules 330, 340, and 350 are said to be context clients because they consume attributes. For example, context client 340 consumes user.location attribute 341. It can be seen that certain software modules may act both as a context server and as a context client. For example, software module 330 is both a context server and a context client, as it provides the user.in_region attribute 331, as well as consuming a user.location attribute 332. It can also be seen that a context client consumes more than one attribute. For example, context client 350 consumes both user.in_region attribute 351 and user.elevation attribute 352. To determine which attributes are currently available, any of the context clients may request that the characterization module enumerate the available attributes. In response to such a request, the characterization module would enumerate the user.location attribute, the user.location attribute, and the user.in_region attribute. Additional preferred embodiments permit context clients to otherwise limit attribute enumeration requests. For example, a context client may request enumeration of all user interface attributes by specifying the wildcard-expanded attribute name platform.user_interface.*.

The characterization module preferably implements the API functions described in detail below in order to obtain attribute values from context servers and provide them to context clients. While these functions are preferably exposed via Microsoft Component Object Module ("COM") interfaces, those skilled in the art will

appreciate that they could be made available to context servers and context clients through a variety of other mechanisms, including a number of different function-calling and message-passing invocation schemes.

LaunchContextServer

Any application can launch a context server that is not already running by calling this function, passing the following parameter:

- Context Server Identifier – A filename, COM object, or other identification method that unambiguously identifies the context server and exposes a mechanism for launching it.

This function returns an error if the requested context server is not found, and returns an error if the requested context server is already running.

Each context server has an associated startup behavior, which specifies whether or not the context server is automatically launched when the characterization module launches. Context servers installation application or the characterization module's control panel applet may be used to change the startup behavior.

RegisterContextServer

When a context server launches, it calls this function to notify the characterization module of its intent to provide attribute instances, passing the following parameters:

- Context Server Name – A name that conforms to attribute naming specifications and uniquely identifies the context server.
- Version – A floating point number that corresponds to the version of the context server.
- Installation Date – The date and time that the context server was installed.

- Filename – The context server's main filename expressed as an absolute pathname.
- Request Handler – A reference to a context server interface that accepts messages from the characterization module to the context server. This information is not available to external modules.

If the registration request includes a context server name that is already in use, the characterization module returns an error and does not register the context server.

Figure 4 is a data structure diagram showing a context server table that maintains a portion of the state of the characterization module. When the RegisterContextServer function is called, the characterization module preferably adds a row to the context server table corresponding to the registered context server. The context server table 400 contains rows 401-403, each corresponding to a registered context server. Each row contains a context server name field 411 containing the name of the context server, a version field 412 identifying the version of the context server, an installation date 413 identifying the date on which the context server was installed on the computer system, a filename 414 identifying a primary file in the file system representing the context server, and a request handler field 415 containing a reference to a request handler function on the context server that may be called by the characterization module to request evaluation of one or all of the attributes provided by the context server. The characterization module preferably adds a row to the context server table when the LaunchContextServer function is called. For example, when the location_region_analysis context server calls the LaunchContextServer function, the characterization module preferably adds row 403 to the context server table. The contents of row 403 indicate that version 1.00.315 of the location_region_analysis context server, installed on 2/10/2000 and having file name l_r_a.exe is registered with the request handler referenced in the request handler field 415 of row 403.

CreateAttributeInstance

Context servers may create an attribute instance at any time by calling this function, passing the following parameters:

- Context Server Name – A name that is unique to the requesting context server. This name should be the same for all attributes that the context server creates. In some embodiments, this parameter is implicit in each call to this function, and need not be explicitly passed.
- Attribute Name – The name of the attribute.
- Data Type – The type of data in which the attribute's value and uncertainty are expressed.
- Format Version – The number of the format version in which the value is expressed.
- Request Handler – A reference to the context server function that processes attribute requests from the characterization module. The characterization module may also send messages to the context server via this function. A single request handler may be used for multiple attributes. In an alternative embodiment, the Request Handler parameter may be specified only during context server registration, rather than during attribute instance creation, or during either activity.
- Startup Behavior – Specifies whether or not the context server should be loaded automatically at startup. This parameter is optional. If included, it supercedes any previous setting for the requesting context server. If the startup behavior is never specified, the characterization module is not responsible for launching the context server. In an alternative embodiment, the Startup Behavior parameter may be specified only during context server registration, rather than during attribute instance creation, or during either activity.

This function returns an error if the name matches that of an attribute that the context server has already created. If another instance of the attribute already exists

and has the same format version, this function returns an error if the data type does not match that of the existing attribute.

Figure 5 is a data structure diagram showing an attribute instance table that maintains a portion of the state of the characterization module. The attribute instance table contains a row for each attribute instance created by a context server. When a context server calls the CreateAttributeInstance function, the characterization module preferably adds a row to the attribute instance table. The attribute instance table contains rows 501-504, each corresponding to a different attribute instance. Each of these rows contains the following fields: an attribute name field 511 containing the name of the attribute, a context server name field 512 identifying the context server that created the attribute instance, a value field 513 containing the value of the attribute last provided by the context server, and uncertainty field 514 identifying the level of uncertainty of the value, a timestamp 515 indicating the time at which the value is effective, and a units field 516 identifying the units for the value and the uncertainty. For example, row 501 indicates that an instance of the user.location attribute from the gps context server has the effective time of 13:11:04.023 on 2/22/2000. Its value is 47° 36.73' N, 122° 18.43' W degrees and minutes, with an uncertainty of 0° .09'. It should be noted, as shown, the user's location, expressed in terms of latitude and longitude, is maintained in a single attribute. In alternative embodiments, such a location could be maintained in a larger number of attributes. For example, rather than being maintained in a single user.location attribute, such a location could be distributed across four separate attributes: user.location.latitude.degrees, user.location.latitude.minutes, user.location.longitude.degrees, and user.location.longitude.minutes. The characterization module preferably adds row 504 to the attribute instance table when the location_region_analysis context server calls the CreateAttributeInstance function for the user.in_region attribute.

RegisterContextClient

Each context client registers with the characterization module, passing the following parameters:

- Name – A name for the context client used for all attribute registrations for a particular context client. If the context client also acts as a context server, it preferably uses the identical name for context server registration, though may use a different name.
- Message Handler – A reference to a context client function that processes messages from the characterization module, such as notifications that attributes are destroyed or the characterization module is shutting down.

The first time a context client registers, it provides a non-null message handler. Subsequent calls by that context client may provide a new message handler that replaces the previous one.

Figure 6 is a data structure diagram showing a context client table that maintains a portion of the state of the characterization module. The context client table 600 contains rows 601-603, each corresponding to a registered context client. Each row contains a context client name field 611 identifying the name of the registered context client, as well as a message handler field 612 containing the reference to a message handler provided by the context client for processing messages from the characterization module. The characterization module preferably adds a row to the context client table for each context client that calls the RegisterContextClient function. For example, the characterization module preferably adds row 603 when the region_analysis context client calls the RegisterContextClient function. Added row 603 contains a reference to the message handler for the region_analysis context client.

EnumerateAttributes

Context clients call this function to request a listing of attributes that meet certain criteria, passing the following parameter:

- Source – The name of the context server providing the attributes that the context client is requesting. If supplied, the characterization

module returns all attributes from the corresponding context server.
If not supplied, the characterization module returns all attributes.

This function returns the following information to the calling context client:

- Attributes – A list of the attributes that meet the request's criteria. If no attributes meet the criteria, the list is empty. If an attribute meets the criteria but its source prohibits the context client from seeing the attribute, it is not included in the list.

If a source is specified and not found, the characterization module returns an error to the context client. In additional preferred embodiments, other parameters may be included to limit the set of attributes that are enumerated.

RegisterToConsumeAttributeOrInstance

A context server calls this function for each attribute or attribute instance that it will be requesting, passing the following parameters:

- ContextClientName – The context client's name
- AttributeName – The name of the attribute that the context client is requesting to consume.
- Source – The context server whose instance of this attribute the context client is registering to consume. If blank, the request constitutes a request to consume a mediated attribute, rather than an attribute instance.

Context clients may register for attributes at any time. Requests for nonexistent attributes produce an error. Requests are not persistent, so the context client must register again to consume the attribute or attribute instance when it becomes available.

Figure 7 is a data structure diagram showing an attribute or instance registration table that maintains a portion of the state of the characterization module. The attribute or instance registration table contains rows 701-704, each corresponding to a different registration that a context client has made for an attribute or an attribute instance. Each row contains a context client name field 711 indicating the name of the context client registering to consume the attribute or attribute instance, an attribute name field containing the attribute name for the attribute or attribute instance, and a context server name field 713 that, for attribute instance registrations contains the name of the context server providing the instance, which is blank for attribute registrations. For example, row 704 indicates that the location_map context client has registered to consume the instance of the user.in_region provided by the region_analysis context server.

GetAttribute

Context clients call this function to request an attribute value, passing the following parameters:

- Name – A name identifying the attribute.
- Source – The name of the context server from which to obtain the requested attribute. If no Source is specified, the characterization module assumes the attribute may come from any source, and it uses an attribute mediator to select one.
- Attribute mediator – The name of the selection method that it would like used to select the attribute from multiple instances. If no attribute mediator is provided, the characterization module's default method is used. If a source and an attribute mediator are provided, the mediator is ignored. The characterization module preferably provides attribute mediators such as the following: a mediator that selects the first attribute instance that was created, a mediator that selects the last attribute instance that was created, a mediator that selects the first attribute instance to be provided in response to a

request for evaluation, a mediator that selects the attribute instance having the lowest level of uncertainty, a mediator that selects the attribute instance having the newest data value, a mediator that calculates the average of the available attribute instances, a mediator that selects the most common of the attribute instances, a mediator that allows the user to select among the attribute instances, and mediators that select an attribute instance whose confidence level, as adjusted for attenuation over the age of the data at a variety of rates, is the highest. Those skilled in the art will appreciate that additional types of attribute mediators may also be employed by the characterization module and requested by context clients.

- **Age** – An expression of the maximum age of the attribute information. If the request is made for a specific attribute instance by specifying a source, then only that instance is checked, and if too old, it is freshened. If no source is specified and multiple instances are present, the characterization module applies the attribute mediator to those that satisfy the age requirement. If no instances satisfy the age requirement, all instances are freshened prior to the application of the attribute mediator.
- **Timeout** – A period of time within which the characterization module should return the attribute. If the attribute is not evaluated within this time, the characterization module sends an error message to the context client. The resulting return values are those of the most recent valid evaluation, or null if no valid data are available. A timeout value of zero instructs the characterization module to wait indefinitely for a response.
- **Supplemental Properties** – The context server may also return to the characterization module names and values for supplemental properties associated with the attribute instance. For example, some supplemental properties may include security information used to determine which context clients may receive attribute instance, as

well as additional information further characterizing the attribute instance.

If the request timeout is exceeded and the characterization module subsequently responds to the request, the returned value replaces the current most recent value cached in the characterization module. If an error is returned after the timeout period, it is ignored. The characterization module preferably does not send additional information to the context client that initiated the request.

This function returns the following to the calling context client:

- Name – A string that identifies the attribute.
- Value – The quantity of the attribute.
- Uncertainty – A quantity that represents the range of likely values around Value that the attribute is likely to have. The contents of the Uncertainty property depend upon what type of information is stored in Value. The Uncertainty property characterizes the extent to which the attribute value may differ from the real quantity being measured, based upon factors including -- but not limited to -- the precision of measuring apparatus used to measure the quantity, the conversion of continuous quantities to discrete values in the measurement process, random fluctuations, temporal variations, systematic errors, and measurement latency.
- Timestamp – A timestamp that represents the effective age of the data. In particular, the timestamp specifies the time at which the attribute value would have been valid had it been measured directly at that moment, irrespective of whether the attribute value was actually measured at that time.
- Units – The units of measure in which the Value and Uncertainty have been expressed.
- Source – A reference to the object that created and owns the instance of the attribute, usually a context server. If the attribute comes from a condition or a condition monitor, the source is the source of the

condition or condition monitor. Some attributes that refer to the characterization module or that are monitored by the characterization module have the characterization module as the source.

- Attribute mediator – The name of the selection method that was used to select the attribute from multiple instances. If no attribute mediator was used (either because only one instance of the attribute was available or because a source was specified), the attribute mediator is left blank.
- Format Version – The version of the attribute’s format specification.
- Flags – Not in use at this time.
- Supplemental Properties – Where the context server has provided supplemental properties, they are returned to the context client, except where the supplemental properties are designated to not be provided to context clients.
- Reference to Context Server’s Callback Function – Optional: The characterization module may return to the context client a direct reference to the context server that provided the returned attribute instance, enabling the context client to call the context server directly in the future to get a new value for the attribute.

This function returns an error to the calling context client if the requested attribute, requested source, or the requested mediator does not exist or if its value is unavailable. Attributes whose values are unavailable are ignored if other instances of the attribute satisfy the request made by the context client.

Figure 8 is a flow diagram showing the steps preferably performed by the characterization module when the GetAttribute function is called. In step 801, if the requested attribute exists, then the facility continues in step 803, else the facility continues in step 802 to return an “attribute not found” error. In step 803, if a single instance of the attribute was requested, then the facility continues in step 804, else the facility continues in step 811. In step 804, if the requested instance exists, then the facility continues in step 806, else the facility continues in step 805 to return an “attribute instance not found” error. In step 806, if the age criterion specified for the attribute

request is satisfied, then the facility continues in step 807 to return the requested attribute instance, else the facility continues in step 808 to “freshen” the attribute instance by calling the appropriate context servers’ request handler to request evaluation of the attribute instance. In step 809, if the age criterion is satisfied by the freshened attribute instance, then the facility continues in step 807 to return the freshened attribute instance, else the facility continues in step 810 to return the freshened attribute instance with an “age not satisfied” error.

In step 811, where a single attribute instance was not requested, if any registered instances of the attribute satisfy the age criterion, then the facility continues in step 816, else the facility continues in step 812. In step 812, the facility freshens all registered instances of the requested attribute. In step 813, if any of the attributes freshened in step 812 satisfy the age criterion, then the facility continues in step 816, else the facility continues in step 814. In step 814, the facility applies the requested attribute mediator to select one instance, or otherwise derive a value from the registered instances. In step 815, the facility returns the instance with an “age not satisfied” error.

In step 816, where one or more instances satisfy the age criterion, if more than one instance satisfies the age criterion, then the facility continues in step 817, else the facility continues in step 807 to return the attribute instance that satisfies the age criterion. In step 817, the facility applies the requested attribute mediator to select one instance from among the instances that satisfy the age criterion, or to otherwise derive a value from the instances that satisfy the age criterion. After step 817, the facility continues in step 807 to return the value produced by the mediator.

The characterization module requests attributes from a context server by calling the request handler registered by the context server, passing the following parameters:

- Attribute Name – The name of the attribute being requested.
- Timeout – The timeout period that the context server is expected to fulfill.

The context server is expected to satisfy the request within the timeout period with the following information:

- Value – The quantity of the attribute.
- Uncertainty – A quantity that represents the range of likely values around Value that the attribute is likely to have. The contents of the uncertainty property depend upon what type of information is stored in Value. The uncertainty is required and is expressed in the same units and format as the value.
- Timestamp – A timestamp that represents the effective age of the data. This age reflects the point in time when the value is most accurate. It is the responsibility of the context server to determine how the timestamp relates to those of attributes that it uses. The timestamp is required, and is preferably determined with the computing platform's clock to facilitate easy comparison in context clients.
- Units – The units of measure in which the Value and Uncertainty have been expressed as a predefined set of allowable units, from among a predetermined set of allowable units. The units parameter is required.
- Format Version – The version of the format specification used to express the attribute's value and uncertainty.

If the context server is able to determine that the request cannot be answered, it returns an error to the characterization module. For instance, the context server may be asking a physical sensor for data, and the physical sensor is not responding.

If the context server cannot provide an attribute because it received an error message from an attribute that it queries, it returns the error message that it received.

When a call to this function necessitates the reevaluation of an attribute instance, the characterization module preferably substitutes that value in the corresponding row of the attribute instance table, thereby caching the attribute instance for possible future use. Figure 9 is a data structure diagram showing updated contents of the attribute instance table. It can be seen in attribute instance table 900 that, upon

reevaluation by the ips context server of its instance of the user.elevation attribute, the characterization module replaced the former contents of the value, uncertainty and timestamp field of row 903, shown in Figure 7, with the values resulting from the reevaluation shown in Figure 9.

5 GetAllAttributeInstances

Context clients call this function to request all instances of a particular attribute. This function uses the same parameters as GetAttribute, except that the attribute source is omitted. This function returns the same information for each attribute instance.

In one embodiment, this function is separate from the GetAttribute function. In an alternative embodiment, this function is omitted and the GetAttribute function is called without specifying the attribute source parameter in order to request all instances of a particular attribute.

CompleteContextServerEvaluation

To force a context server to calculate all of its output attributes at the same time, a context client calls this function, passing the following parameters:

- Context Server Name – An identifier of the context server of interest.
- Timeout – A period of time within which the characterization module should successfully force a complete evaluation. If the context server's attributes are not completely evaluated within this time, the characterization module sends an error message to the context client.

When this function is called, it in turn calls the request handler for the identified context server, requesting complete evaluation of the context server's attributes by specifying a special value for the AttributeName parameter. In response, the context

server evaluates all of its attributes and provides them to the characterization module delaying other requests until it has finished. If the context server uses other attributes to evaluate its own, it requests fresh versions of these attributes, but it is neither required to or prohibited from requesting complete evaluation.

This function returns the following to the calling context client:

- Attributes – The attributes that were simultaneously evaluated including all of their constituent information as specified in the GetAttribute section. Note that the attribute source need not be repeated for each attribute.

This function returns an error if the request either timed-out or there was no context server with the specified name.

The characterization module possesses several attributes of its own that may be of interest to context clients and context servers. Read-only parameters are surfaced as attributes identical to those created by other context servers, except that the owner of these attributes is the characterization module itself. Several adjustable items require two-way information flow and thus have their own interfaces. In an alternative embodiment, adjustable items are exposed as attributes.

GetCharacterizationModuleAttribute

Context clients call this function to request attributes containing information about the characterization module, passing a parameter identifying the requested attribute. The available items are attributes for which the characterization module is the creator and owner; otherwise they behave identically to all other attributes. characterization module attributes are visible to all context clients. The following read-only characterization module attributes may be requested:

- Version – Contains a description of the characterization module version.

- Path – A string containing the absolute path within which the characterization module has been installed, in the proper format for the computing platform.
- Time – The local time as maintained by the computer's clock.
- Date – The date as maintained by the computer's clock.
- Default Timeout – Some requests require polling context servers for attribute information. The characterization module sends an error to the requesting context client if the context server does not respond within a specified time interval. If the context client has not specified this interval with its request, the characterization module's default value is used. A timeout value of zero instructs the characterization module to wait indefinitely for a response. No parameters are necessary to read the default timeout setting, and such requests result in an immediate response from the characterization module with that information. The default timeout may preferably be changed by the user via a characterization module control application/control panel applet. In additional embodiments, the default attribute mediator may be changed in another manner, such as under the control of a program other than the characterization module.
- Default Attribute Mediator – Context clients can read the characterization module's default attribute mediator. No parameters are necessary to read the default attribute mediator, and the characterization module returns the attribute mediator's name. The default attribute mediator may preferably be changed by the user via a characterization module control application/ control panel applet. In additional embodiments, the default attribute mediator may be changed in another manner, such as under the control of a program other than the characterization module.

In one embodiment, this function is separate from the GetAttribute function. In an alternative embodiment, this function is omitted and the GetAttribute function is called in order to request a characterization module attribute.

CreateCondition

Context clients call this function to create a condition to check a relationship between an attribute and a value or between two attributes, passing the following:

- Name – A name for the condition. This name must be unique among all conditions created by the same context client.
- 1st Logical Parameter – An attribute name identifying an attribute, an attribute name and source identifying an attribute instance, or a condition name identifying a condition.
- 2nd Logical Parameter – An attribute name identifying an attribute, an attribute name and source identifying an attribute instance, or a condition name identifying a condition.
- Value – A value for attribute comparison. If an attribute is provided as the 2nd logical parameter, the value is ignored.
- Logical Operator – One of a predefined set of logical operators. The allowed operators depend upon whether the characterization module is asked to compare attributes or conditions as shown in Table 1 below.

Attributes	Conditions
>	AND
<	OR
=	XOR
>=	NOR
<=	NOT
<>	

Table 1. List of Logical Operators for Conditions

Conditions may compare an attribute with another attribute, and attribute with a value, or a condition with another condition. These combinations are listed below in Table 2.

1 st Parameter	2 nd Parameter	Action
Attribute Name	<none>	Compares attribute to value
Attribute Name	Attribute Name	Compares attributes to each other
Condition Name	Condition Name	Compares conditions to each other

Table 2. Parameter Usage for Conditions

This function returns an error if the requested name is already in use by another condition that was created by the calling context client. An error is returned if the referred to attributes or conditions do not exist. An error is returned if an attribute is requested along with a condition or vice-versa.

Figure 10 is a data structure diagram showing a condition table that contains a portion of the state of the characterization module. Condition table 1000 has a row for each condition created by a context client. Row 1001 contains a condition name field 1011 containing the name of the condition, a context client name 1012, identifying the context client that created the condition, a first logical parameter field 1013 and a second logical parameter field 1014 identifying attributes or other conditions that are compared by the condition, a comparison value 1015 that specifies a value to which the first logical parameter is compared if no second logical parameter is listed, and a logical operator 1016, identifying the logical operator to be applied in the comparison. The

characterization module preferably adds row 1001 to condition table 1000 when the region_analysis context client creates the in_region_true condition to indicate whether the user is presently within a special region. Added row 1001 indicates that the in_region_true condition was created by the region_analysis context client, and has first
5 logical parameter user.in_region, no second logical parameter, comparison value TRUE, and logical operator “=”.

EvaluateCondition

A context client calls this function to query conditions that it has created passing the following parameters:

- Name – The name of the condition.
- Timeout – A period of time within which the characterization module should have successfully evaluated the condition.

This function returns the following to the calling context client:

- Value – A Boolean expression resulting from the evaluation of the condition.

This function returns an error if the condition requests a non-existent attribute, the requested condition does not exist, or the timeout is exceeded.

When this function is called for a particular condition, the facility preferably uses the information in the corresponding row of the condition table to evaluate the condition. For example, when this function is called for the in_region_true condition, the facility preferably uses information row 1001 of condition table 1000
25 shown in Figure 10 to evaluate this condition.

CreateConditionMonitor

A context client calls this function to create a condition monitor to monitor for a change in a condition and be notified when it occurs, passing the following parameters:

- Name – The name of the condition monitor.
- Condition – The name of the condition that triggers the condition monitor.
- Behavior – Specifies when the condition monitor triggers. The condition monitor may be triggered when the condition becomes true, becomes false, or when it changes state in either direction.
- Frequency – The time between subsequent checks for triggering of the condition monitor. This value must be non-zero.
- Trigger Handler – A reference to a context client procedure that the characterization module notifies when the condition monitor is triggered.

An error is returned if the name is already in use by another condition monitor from the calling context client.

Figure 11 is a data structure diagram showing a condition monitor table that maintains a portion of the state of the characterization module. Condition monitor table 1100 has a row 1101 corresponding to a condition and containing each of the following fields: a condition monitor name field 1111 containing the name of the condition monitor; a context client name field 1112 containing the name of the context client that created the condition monitor; a condition name field 1113 that contains the name of the condition monitored by the condition monitor; a behavior field 1114 that indicates whether the condition monitor is triggered when the condition becomes true, when it becomes false, or when it changes value in either direction; a frequency field 1115 showing the frequency with which the condition monitor evaluates the condition; a condition last evaluated field 1116 showing the time at which the condition monitor last evaluated the condition; a trigger handler reference 1117 that identifies a trigger handler

function of the context client that is to be called when the condition monitor is triggered; and a stopped field 1118 that indicates whether the context client has suspended operation of the condition monitor. The condition monitor preferably adds row 1101 to the condition monitor table when the region_analysis context client creates the region_boundary_crossed condition to notify the region_analysis context client when the value of the in_region_true condition changes, indicating that the user has entered or left the special region. Row 1101 indicates that the region_boundary_crossed condition monitor was created by the region_analysis context client, is based upon the in_region_true condition, monitors for both TRUE and FALSE behavior, is evaluated every 30 seconds, was last evaluated at 13:11:29.101 on 2/22/2000, has the trigger handler referenced in field 1117 of row 1101, and is not presently stopped.

StopConditionMonitor

A context client calls this function to temporarily suspend the operation of a condition monitor, passing the following parameter:

- Name – The name of the condition monitor.

This function returns an error to the calling context client if the name does not correspond to an existing condition monitor created by that context client.

Figure 12 is a data structure diagram showing updated contents of the condition monitor table. It can be seen from stopped field 1218 of row 1201 in condition monitor table 1200 that the region_analysis context client has stopped, or suspended the operation of, the region_boundary_crossed condition monitor, perhaps in response to the observation that the user is now asleep and his or her location will remain constant.

StartConditionMonitor

A context client calls this function to resume operation of a stopped condition monitor, passing the following parameter:

- Name – The name of the condition monitor.

This function returns an error to the calling context client if the name does not correspond to an existing condition monitor created by that context client.

5 When the StartConditionMonitor function is called, the facility preferably changes the contents of the stopped field 1118 from “yes” as shown in Figure 11 to “no” as shown in Figure 12, so that the characterization module resumes operation of the condition monitor.

RemoveConditionMonitor

10 Context clients call this function to remove a condition monitor that it has created, passing the following parameter:

- Name – The name of the condition monitor.

15 This function returns an error to the calling context client if the name does not correspond to an existing condition monitor created by the calling context client. If the condition monitor is active when this function is called, this function automatically stops the condition monitor prior to removing it.

20 When this function is called, the characterization module preferably deletes the row of the condition monitor table corresponding to the condition monitor. For example, if this function is called for the region_boundary_crossed condition monitor, the characterization module preferably deletes row 1201 of condition monitor table 1200 shown in Figure 12.

RemoveCondition

25 A context client calls this function to remove a condition that it owns, passing the following parameter:

- Name – the name of the condition to be removed.

An error is returned if the name does not correspond to an existing condition created by the calling context client. The characterization module returns an error if other conditions or condition monitors rely upon the condition, and the condition is not removed.

When this function is called for a particular condition, the characterization module preferably deletes the corresponding row from the condition table. For example, when this function is called for `in_region_true` condition the facility preferably deletes row 1001 from condition table 1000 shown in Figure 10.

In the foregoing, a preferred embodiment is described in which conditions and condition monitors are created separately. In an alternative preferred embodiment, conditions and condition monitors are created together as a single entity, with a single API function call.

UnregisterToConsumeAttributeOrInstance

A context server calls this function to unregister to consume an attribute or attribute instance that it is registered to consume, passing the following parameters:

- Context Server Name – The name of the context server.
- Attribute Name – The name of the attribute for which the context client is unregistering.
- Source – For requests to unregister to consume an attribute instance, contains the name of the context server providing the attribute instance. For requests to unregister to consume an attribute, is blank.

This function removes the row corresponding to the specified attribute or instance registration from the attribute or instance registration table.

This function returns an error if the attribute or instance registration table does not contain a row corresponding to the specified registration.

UnregisterContextClient

Context clients may unregister themselves at any time. Before calling this function, context clients are expected to first remove all conditions, condition monitors, and attributes that pertain to the registration they are ending, parsing the following:

5

- Context Client Name – Name of the context client to unregister.

Unregistration of a context client results in the removal of all remaining conditions, condition monitors, and attribute consumption registrations that it owns.

10

The characterization module may ask the context client to unregister. Such requests require no parameters. The context client responds to such a request by calling this function within a reasonable period of time. If such acknowledgement is not made, the characterization module preferably removes the context client registration and all of its conditions and condition monitors automatically.

15

When this function is called for a particular context client, the characterization module preferably deletes the corresponding row from the context client table. For example, when this function is called for the region_analysis context client, the facility preferably deletes row 603 from context client table 600 shown in Figure 6.

CheckAttributeInstanceDependencies

20

To determine whether there are context clients using one of its attribute instances, a context server calls this function, passing the following parameter:

- Attribute – The name of the attribute.
- Context Server Name – The name of the calling context server.

25

This function returns the following to the calling context server:

- Number – The number of context clients that are registered for the attribute instance named in the request. This includes all context

clients that have registered for the context server's specific instance as well as those that have registered for the attribute for which there is only one instance. This does not include registrations for an attribute for which there are additional context server's able to satisfy the context clients' requests.

This function returns an error if the requested attribute instance does not exist.

When this function is called, it preferably first refers to the attribute or instance registration table to count how many context clients have registered specifically for the calling context server's instance of the attribute. The function further identifies any registrations for this attribute from any context server, *i.e.*, rows of the attribute or instance registration table containing the attribute name of this attribute instance and no context server name. For these identified rows, the function consults the attribute instance table to determine whether instances of this attribute other than the one provided by the calling context server are present in the table. If not, the function also counts the number of context clients registered to consume the attribute from any context server.

RemoveAttributeInstance

Context servers call this function to remove the instances of attributes that they have created, passing the following parameter:

- Name – The name of the attribute to remove.

This function returns an error if the requested name does not correspond to an attribute that the context server may remove.

This function notifies context clients registered for this attribute instance that it is no longer available. Context clients that register for any one of multiple instances of an attribute are only notified when the last instance of the attribute is removed. If conditions or condition monitors rely upon the attributes that are being removed, their owners are notified, and the conditions and rules are removed.

When this function is called for a particular attribute instance, the characterization module preferably removes the corresponding row from the attribute instance table. For example, when the location_region_analysis context server calls this function for its instance of the user.in_region attribute, the characterization module preferably deletes row 504 from attribute instance table 500 shown in Figure 5.

UnregisterContextServer

A context server may unregister at any time by calling this function. Calling this function requires no additional parameters.

Unregistration causes the characterization module to remove any remaining attribute instances belonging to that context server (thereby requiring the characterization module to notify any registered context clients) prior to actually releasing the context server.

The characterization module may ask the context server to unregister, in which case the context server is expected to acknowledge the request by requesting unregistration. Failure of the context server to acknowledge the request within a reasonable period of time causes the context server and its attributes to be removed automatically.

The user is preferably able to control several values that affect characterization module function using a windows control panel applet dedicated to the characterization module. The parameters under user control include the following:

- Default Timeout – See Default Timeout section.
- Default Attribute Mediator – See Default Attribute Mediator section.
- Context Server Listing – A list of all currently registered context servers and whether or not each is automatically started at characterization module startup.
- Context Client Listing – A list of all currently registered context clients and the attributes for which they have registered.

Users can enter a default timeout value. This value must be greater than or equal to zero.

Users can select the default attribute mediator from a list of available mediators.

5 Users can refresh the display of context servers and context clients at any time.

When UnregisterContextServer function is called by a particular context server, the characterization module preferably removes the corresponding row from the context server table. For example, when this function is called for the
10 location_region_analysis context server, the characterization module preferably removes row 403 from context server table 400 shown in Figure 4.

In the foregoing, the facility is described as being implemented using a characterization module that is called by context servers and context clients, caches
15 attribute values, and maintains status information about the operation of context servers and context clients. In an alternative preferred embodiment, however, the facility operates without the use of such as characterization module. In this embodiment, context servers communicate directly with context clients.

Figure 13 is a data flow diagram showing the operation of the facility
20 without a characterization module. It can be seen in Figure 13 that context servers 1310, 1320, and 1330 provide attributes directly to context clients 1330, 1340, and 1350. For example, it can be seen that context server 1320 provides a user.elevation attribute 1322 directly to context client 1350. In this embodiment, the context client may itself cache attribute values recently obtained from a context server. Further, in this embodiment,
25 context clients may themselves interrogate context servers for an enumeration of their attributes, and mediate between attribute instances provided by different context servers. For example, context client 1340 may mediate between the instance 1311 of the user.location attribute provided by context server 1310 and the instance 1321 of the user.location attribute provided by context server 1320.

30 Additional embodiments of the facility support supplemental properties for attribute instances that are supplied by the context server supplying an attribute instance,

maintained by the characterization module, and, in some cases, returned to context clients requesting a value of the corresponding attribute.

Figure 14 is a data structure diagram showing an attribute instance property table that maintains a portion of the state of the characterization module. Each row in the attribute instance property table 1400 corresponds to a supplemental property provided by a context server for an attribute instance it provides. Each row contains an attribute name field containing the attribute name for the attribute instance, a context server name 1402 containing the name of the context server providing the attribute instance, a property name 1403 containing a property name for the property, and a property value field 1404 containing the value of the property. For example, row 1401 shows that with its instance of the user.location attribute, the ips context server has also provided a security_token property having the value 5A637AR. Some supplemental properties are preferably provided to context clients calling the GetAttribute function for the attribute instance. Other supplemental properties, such as the security_token property, represented by row 1401, are preferably withheld from context clients for the exclusive use of the characterization module. For example, the characterization module preferably uses the security_token property to determine whether a particular context client should be permitted to obtain values of the user_location attribute instance provided by the ips context server by calling the GetAttribute function.

In additional preferred embodiments, the facility may operate with a partial characterization module. Such a partial characterization module may include various combinations of the functionalities of routing communication between context servers and the context clients that consume their attributes, caching attribute values, enumerating available attributes, and providing attribute mediation. In further preferred embodiments, the facility utilizes a characterization module that constitutes a passive data store that is shared between context servers and context clients. Context servers access the data store to write attribute values, and context clients access the data store to read attribute values.

In additional preferred embodiments, the facility may provide a characterization module that implements a “push” information flow model, in which, each time an attribute value provided by a context server changes, the new value is automatically provided to context clients. In some cases, context servers may push selected attribute values. For example, context servers may push attributes whose values

are constant, attributes whose values change infrequently, or attributes whose values are expensive to obtain. In further preferred embodiments, the facility provides a characterization module that implements a “pull” information flow model, in which attribute values are only obtained by the characterization module from the context servers
5 when they are requested by a context client. In additional preferred embodiments, characterization modules are provided that support a variety of other information flow models.

It will be understood by those skilled in the art that the above-described facility could be adapted or extended in various ways. For example, the characterization
10 module may be implemented in computing devices other than wearable computers, including other types of mobile computing devices, such as personal digital assistants. The characterization module may further be implemented on various types of stationary and/or distributed computing devices, as well as non-traditional computing devices, such as smart appliances. Further, rather than the attribute, context server, context client,
15 condition, and condition monitor names discussed above, the facility may use other types of identifiers, such as handles. While the foregoing description makes reference to preferred embodiments, the scope of the invention is defined solely by the claims that follow and the elements recited therein.

CLAIMS

We claim:

1. A method in a computing device for exchanging context attributes,
1 comprising:

3 receiving an invocation request to provide an attribute value, the invocation
4 request being generated by a requesting attribute consumer, the invocation request
5 identifying the attribute whose value is to be provided; and

6 in response to receiving the invocation request, providing a value for the
7 identified attribute to the requesting attribute consumer.

2. The method of claim 1 wherein the invocation request is a function call.

3. The method of claim 1 wherein the invocation request is a procedure.

4. The method of claim 1 wherein the invocation request is an invocation
1 message.

5. The method of claim 1 wherein a value of the identified attribute is
1 stored, and wherein the stored value is provided to the requesting attribute consumer.

6. The method of claim 1 wherein the identified attribute is associated with
1 an attribute source, and wherein the method further comprises obtaining a value of the
2 identified attribute from the attribute source with which the identified attribute is associated,
3 and wherein the value of the attribute obtained from the attribute source with which the
4 identified attribute is provided to the requesting attribute consumer.

7. The method of claim 1 further comprising, in addition to providing a
1 value for the identified attribute to the requesting attribute consumer, providing units of the
2 value for the identified attribute.

8. The method of claim 1 further comprising, in addition to providing a
1 value for the identified attribute to the requesting attribute consumer, providing an
2 uncertainty level for the identified attribute.

9. The method of claim 1 further comprising, in addition to providing a
1 value for the identified attribute to the requesting attribute consumer, providing a timestamp
2 for the identified attribute.

10. The method of claim 1 wherein the identified attribute is information
1 reflecting an aspect of the computing device.

11. The method of claim 10 wherein the computing device has a visual
1 output device, and wherein the identified attribute is information about the availability of the
2 visual output device.

12. The method of claim 1 wherein the computing device is present in an
1 environment, and wherein the identified attribute is information reflecting an aspect of the
2 environment.

13. The method of claim 12 wherein the environment has a temperature, and
1 wherein the identified attribute is the temperature of the environment.

14. The method of claim 1 wherein the computing device has a user, and
1 wherein the identified attribute is information reflecting an aspect of the user.

15. The method of claim 14 wherein the user has a blood pressure, wherein
1 the identified attribute is the blood pressure of the user.

16. The method of claim 1 wherein one or more applications are executing
1 on the computing device, and wherein the identified attribute is information reflecting an
2 aspect of an executing application.

17. The method of claim 16 wherein an electronic messaging application is
1 among the applications executing on the computing device, and wherein the identified
2 attribute indicates whether new messages have been received by the electronic messaging
3 application.

18. The method of claim 1 wherein the computing device is outside a
1 selected remote environment, and wherein the identified attribute is information reflecting an
2 aspect of the remote environment.

19. The method of claim 18 wherein the environment has a temperature, and
1 wherein the identified attribute is the temperature of the remote environment.

20. The method of claim 1 wherein the computing device has a user, and
1 wherein the identified attribute is information reflecting an aspect of a selected person other
2 than the user.

21. The method of claim 20 wherein the selected person has a location, and
1 wherein the identified attribute is the location of the selected person.

22. The method of claim 1 wherein the identified attribute is information
1 reflecting an aspect of an identified person.

23. The method of claim 22 wherein the identified person has a temperature,
1 and wherein the identified attribute is the temperature of the identified person.

24. The method of claim 1 wherein the invocation request further specifies a
1 maximum age for the attribute value, and wherein the value provided to the requesting
2 attribute consumer has an age that is no older than the specified maximum age.

25. The method of claim 1 wherein the identified attribute has a source, and
1 wherein a value of the identified attribute previously obtained from the source of the
2 identified attribute is cached with an indication of the age of the previously-obtained value of
3 the identified attribute, and wherein the method further comprises, in response to receiving
4 the invocation request:

6 determining whether the age of the previously-obtained value of the identified
7 attribute exceeds the specified maximum age;

8 if the age of the previously-obtained value of the identified attribute does not
9 exceed the specified maximum age, providing the previously-obtained value of the identified
10 attribute to the requesting attribute consumer; and

11 if the age of the previously-obtained value of the identified attribute exceeds
12 the specified maximum age:

13 obtaining a new value of the identified attribute from the source of the
14 identified attribute, and

15 providing the new value of the identified attribute to the requesting
16 attribute consumer.

26. The method of claim 34 further comprising caching the new value of the
1 identified attribute.

27. The method of claim 1 wherein the identified attribute has a source, and
1 wherein a value of the identified attribute previously obtained from the source of the
2 identified attribute is cached with an indication of the age of the previously-obtained value of
3 the identified attribute, and wherein the method further comprises, in response to receiving
4 the invocation request:

6 determining whether the age of the previously-obtained value of the identified
7 attribute exceeds the specified maximum age;

8 if the age of the previously-obtained value of the identified attribute does not
9 exceed the specified maximum age, providing the previously-obtained value of the identified
10 attribute to the requesting attribute consumer; and

11 if the age of the previously-obtained value of the identified attribute exceeds
12 the specified maximum age:

13 obtaining a new value of the identified attribute from the source of the
14 identified attribute,

15 determining that the new value of the identified attribute exceeds the
16 specified maximum age, and

17 in response to determining that the new value of the identified attribute
18 exceeds the specified maximum age, generating an error condition indicating that the source
19 is unable to satisfy the specified maximum age.

28. The method of claim 1 wherein a value of the identified attribute is
1 available from each of a plurality of sources, and wherein the invocation request further
2 specifies one of the plurality of sources from which a value of the identified attribute is
3 available, and wherein the value of the identified attribute provided to the requesting attribute
4 consumer is from the specified source.

29. The method of claim 1 wherein a value of the identified attribute is
1 available from each of a plurality of sources, and wherein the invocation request further
2 specifies to provide a value of the identified attribute from each of the sources from which a
3 value of the identified attribute is available, and wherein a value of the identified attribute
4 from each of the sources from which a value of the identified attribute is available is
5 provided to the requesting attribute consumer.

30. The method of claim 1 wherein a value of the identified attribute is
1 available from each of a plurality of sources, and wherein the value of the identified attribute
2 provided to the requesting attribute consumer is determined based upon the values of the
3 identified attribute available from the plurality of sources.

31. The method of claim 30, further comprising caching the value of the
1 identified attribute provided to the requesting attribute consumer.

32. The method of claim 1 wherein a plurality of values of the identified
1 attribute are available from different sources, and wherein the invocation request further
2 specifies to select one of the plurality of available values of the identified attribute to
3 provide, and wherein the value of the identified attribute provided to the requesting attribute
4 consumer is selected from the plurality of available values of the identified attribute.

33. The method of claim 1 wherein a plurality of values of the identified
1 attribute are available from different sources, and wherein the invocation request further
2 specifies to determine a value of the identified attribute to provide that is based upon the
3 plurality of available values of the identified attribute but different than each of the plurality
4 of available values of the identified attribute, and wherein the value of the identified attribute
5 provided to the requesting attribute consumer is based upon the plurality of available values
6 of the identified attribute but different from each of the plurality of available values of the
7 identified attribute.

34. The method of claim 1 wherein a plurality of values of the identified
1 attribute are available from different sources, and wherein the invocation request further
2 specifies to determine, from the plurality of available values of the identified attribute, one
3 value of the identified attribute to provide, and wherein the value of the identified attribute
4 provided to the requesting attribute consumer is determined from the plurality of available
5 values of the identified attribute.

35. The method of claim 34 wherein the invocation request further specifies
1 a basis for determining, from the plurality of available values of the identified attribute, one
2 value of the identified attribute to provide, and wherein the value of the identified attribute
3 provided to the requesting attribute consumer is determined from the plurality of available
4 values of the identified attribute using the specified basis.

36. The method of claim 35 wherein the invocation request further specifies
1 selecting the oldest available value of the identified attribute, and wherein the value of the
2 identified attribute provided to the requesting attribute consumer is the oldest available value
3 of the identified attribute.

37. The method of claim 35 wherein the invocation request further specifies
1 selecting the newest available value of the identified attribute, and wherein the value of the
2 identified attribute provided to the requesting attribute consumer is the newest available
3 value of the identified attribute.

38. The method of claim 35 wherein the available values of the identified
1 attribute are each requested and received from a different source, and wherein the invocation
2 request further specifies selecting the first-received available value of the identified attribute,
3 and wherein the value of the identified attribute provided to the requesting attribute
4 consumer is the first-received available value of the identified attribute.

39. The method of claim 35 wherein each available value of the identified
1 attribute has an uncertainty level, and wherein the invocation request further specifies
2 selecting the available value of the identified attribute having the lowest uncertainty level,
3 and wherein the value of the identified attribute provided to the requesting attribute
4 consumer is the available value of the identified attribute having the lowest uncertainty level.

40. The method of claim 35 wherein the invocation request further specifies
1 determining a value of the identified attribute that is the average of the available values of the
2 identified attribute, and wherein the value of the identified attribute provided to the
3 requesting attribute consumer is the average of the available values of the identified attribute.

41. The method of claim 35 wherein each of the available values of the
1 identified attribute has an uncertainty level, and wherein the value of the identified attribute
2 provided to the requesting attribute consumer is the average of the available values of the
3 identified attribute, weighted by their uncertainty levels.

42. The method of claim 35 wherein each of the available values of the
1 identified attribute has an age, and wherein the value of the identified attribute provided to
2 the requesting attribute consumer is the average of the available values of the identified
3 attribute, weighted by their ages.

43. The method of claim 35 wherein the invocation request further specifies
1 selecting the available value of the identified attribute that occurs the largest number of times
2 among the available values of the identified attribute, and wherein the value of the identified
3 attribute provided to the requesting attribute consumer is the available value of the identified
4 attribute that occurs the largest number of times among the available values of the identified
5 attribute.

44. The method of claim 35 wherein the invocation request further specifies
1 soliciting selection by a user of one of the available values of the identified attribute, and
2 wherein the value of the identified attribute provided to the requesting attribute consumer is
3 an available value of the identified attribute selected by the user.

45. The method of claim 35 wherein each available value of the identified
1 attribute has an uncertainty level and an effective time, and wherein the invocation request
2 further specifies selecting an available value of the identified attribute based upon a function
3 of the uncertainty level and effective time of each, and wherein the value of the identified
4 attribute provided to the requesting attribute consumer is selected from the available values
5 of the identified attribute based upon the function of the uncertainty level and effective time
6 of each.

46. The method of claim 1 wherein the invocation request is received by a
1 context characterization module, and wherein the context characterization module has an
2 attribute under its control, and wherein the invocation request identifies the attribute under
3 the control of the characterization module, and wherein the attribute whose value is provided
4 to the requesting attribute consumer is the attribute under the control of the characterization
5 module.

47. The method of claim 1 further comprising, before the invocation request
1 to provide an attribute value is received, receiving from the attribute consumer an invocation
2 request to register the attribute consumer.

48. The method of claim 47 wherein the invocation request to register an
1 attribute consumer indicates that the requesting attribute consumer reserves an opportunity to
2 later request provision of a value of the identified attribute, the method further comprising
3 associating with the identified attribute an indication that the requesting attribute consumer is
4 dependent on the identified attribute.

49. The method of claim 1 further comprising, before the invocation request
1 to provide an attribute value is received, receiving an invocation request to register an
2 attribute source for the identified attribute, and wherein the method further comprises
3 obtaining a value of the identified attribute from the attribute source with which the
4 identified attribute is associated, and wherein the value of the attribute obtained from the
5 attribute source with which the identified attribute is provided to the requesting attribute
6 consumer.

50. A computing device for exchanging context attributes, comprising:
2 an invocation request receiver that receives an invocation request to provide an
3 attribute value, the invocation request being generated by a requesting attribute consumer,
4 the invocation request identifying the attribute whose value is to be provided; and
5 an attribute value provider that provides a value for the identified attribute to
6 the requesting attribute consumer in response to receipt of the invocation request by the
7 invocation request receiver.

51. The computing device of claim 50 wherein the computing device is a
1 mobile computing device.

52. The computing device of claim 50 wherein the computing device is a
1 wearable computing device worn on the body of a human user.

53. The computing device of claim 50 further comprising an attribute value
1 memory that contains a value of the identified attribute, wherein the attribute value contained
2 in the attribute value memory is provided to the requesting attribute consumer by the
3 attribute value provider.

54. The computing device of claim 50 wherein the identified attribute is
1 associated with an attribute source, the computing device further comprising an attribute
2 procurement subsystem that obtains a value of the identified attribute from the attribute
3 source with which the identified attribute is associated, and wherein the value of the attribute
4 obtained from the attribute source with which the identified attribute is provided to the
5 requesting attribute consumer by the attribute value provider.

55. The computing device of claim 50 wherein the identified attribute has a
1 source, and wherein a value of the identified attribute previously obtained from the source of
2 the identified attribute is cached with an indication of the age of the previously-obtained
3 value of the identified attribute, and wherein the computing device further comprises an age
4 determination subsystem that determines whether the age of the previously-obtained value of
5 the identified attribute exceeds the specified maximum age in response to receipt of the
6 invocation request by the invocation request receiver, and wherein the attribute value
7 provider provides the previously-obtained value of the identified attribute to the requesting
8 attribute consumer if the age of the previously-obtained value of the identified attribute does
9 not exceed the specified maximum age, and obtains a new value of the identified attribute
10 from the source of the identified attribute and provides the new value of the identified
11 attribute to the requesting attribute consumer if the age of the previously-obtained value of
12 the identified attribute exceeds the specified maximum age.

56. The computing device of claim 50 wherein a value of the identified
1 attribute is available from each of a plurality of sources, and wherein the invocation request
2 further specifies one of the plurality of sources from which a value of the identified attribute
3 is available, and wherein the value of the identified attribute provided to the requesting
4 attribute consumer by the attribute value provider is from the specified source.

57. The computing device of claim 50 wherein a value of the identified
1 attribute is available from each of a plurality of sources, and wherein the invocation request
2 further specifies to provide a value of the identified attribute from each of the sources from
3 which a value of the identified attribute is available, and wherein a value of the identified
4 attribute from each of the sources from which a value of the identified attribute is available is
5 provided to the requesting attribute consumer by the attribute value provider.

58. The computing device of claim 50 wherein a value of the identified
1 attribute is available from each of a plurality of sources, and wherein the value of the
2 identified attribute provided to the requesting attribute consumer is determined based upon
3 the plurality of values of the identified attribute available from different sources.

59. The computing device of claim 50 wherein a plurality of values of the
1 identified attribute are available from different sources, and wherein the value of the
2 identified attribute provided to the requesting attribute consumer by the attribute value
3 provider is selected from the plurality of available values of the identified attribute.

60. The computing device of claim 50 wherein a plurality of values of the
1 identified attribute are available from different sources, and wherein the value of the
2 identified attribute provided to the requesting attribute consumer by the attribute value
3 provider is based upon the plurality of available values of the identified attribute but different
4 from each of the plurality of available values of the identified attribute.

61. A computer-readable medium whose contents cause a computing device
1 to exchange context attributes by:

3 receiving an invocation request to provide an attribute value, the invocation
4 request being generated by a requesting attribute consumer, the invocation request
5 identifying the attribute whose value is to be provided; and

6 in response to receiving the invocation request, providing a value for the
7 identified attribute to the requesting attribute consumer.

62. A computer memory containing an invocation data structure generated
1 by a requesting attribute consumer, the data structure representing an invocation request to
2 provide an attribute value, and comprising:
4 an indication that the represented invocation request is a request for the
5 invocation of a function to provide an attribute value;
6 information identifying the attribute whose value is to be provided; and
7 information identifying a program to which the attribute value is to be
8 provided,
9 such that an attribute value provision system receiving the invocation data structure could use
10 its contents to provide an attribute value in accordance with the request it represents.

63. The computer memory of claim 62 wherein the invocation data structure
1 is a function call stack record.

64. The computer memory of claim 62 wherein the invocation data structure
1 is a function invocation message.

65. A method in a computing device for exchanging context attributes,
1 comprising:
3 receiving an invocation request to provide a list of attributes available in the
4 computing device, the invocation request being generated by a requesting attribute consumer;
5 and
6 in response to receiving the invocation request, providing to the requesting
7 attribute consumer a list of available attributes.

66. The method of claim 65 wherein the attributes available in the
1 computing device are each associated with an attribute source, and wherein the invocation
2 request further specifies an attribute source, and wherein the list of available attributes
3 provided to the requesting attribute consumer contains only those available attributes
4 associated with the specified attribute source.

67. The method of claim 65 wherein each attribute available in the
1 computing device has a name, and wherein the invocation request further specifies an
2 attribute name pattern, and wherein the list of available attributes provided to the requesting
3 attribute consumer contains only those available attributes whose names match the specified
4 pattern.

68. A computer-readable medium whose contents cause a computing device
1 to exchange context attributes by:

3 receiving an invocation request to provide a list of attributes available in the
4 computing device, the invocation request being generated by a requesting attribute consumer;
5 and

6 in response to receiving the invocation request, providing to the requesting
7 attribute consumer a list of available attributes.

69. The computer-readable medium of claim 67 wherein the attributes
1 available in the computing device are each associated with an attribute source, and wherein
2 the invocation request further specifies an attribute source, and wherein the list of available
3 attributes provided to the requesting attribute consumer contains only those available
4 attributes associated with the specified attribute source.

70. A computing device for exchanging context attributes, comprising:
2 an invocation request receiver that receives an invocation request to provide a
3 list of attributes available in the computing device, the invocation request being generated by
4 a requesting attribute consumer; and
5 an available attribute list provider that provides to the requesting attribute
6 consumer a list of available attributes in response to receipt of the invocation request by the
7 invocation request receiver.

71. A computer memory containing an invocation data structure generated
1 by a requesting attribute consumer, the data structure representing an invocation request to
2 provide a list of available attributes, and comprising:

4 an indication that the represented invocation request is a request for the
5 invocation of a function to provide a list of available attributes; and
6 information identifying a program to which the list of available attributes is to
7 be provided
8 such that an available attribute listing system receiving the invocation data
9 structure could use its contents to provide a list of available attributes in accordance with the
10 request it represents.

72. The computer memory of claim 71 wherein each available attribute is
1 associated with an attribute source, and wherein the invocation request further comprises
2 information specifying an attribute source,

4 such that an available attribute listing system receiving the invocation data
5 structure could use its contents to provide a list of only those available attributes associated
6 with the specified attribute source.

73. A method in a computing device for exchanging context attributes,
1 comprising:

3 receiving an invocation request to register an attribute, the invocation request
4 being generated by an attribute source, the invocation request identifying the attribute to be
5 registered; and

6 in response to receiving the invocation request, generating an indication that
7 the identified attribute can be obtained from the attribute source.

74. The method of claim 73 wherein the invocation request further specifies
1 a manner of invoking the attribute source to obtain the identified attribute from the attribute
2 source, and wherein the generated indication that the identified attribute can be obtained
3 from the attribute source includes an indication of the specified manner of invoking the
4 attribute source to obtain the identified attribute from the attribute source.

75. The method of claim 74, further comprising using the indication of the
1 specified manner of invoking the attribute source included in the indication to invoke the
2 attribute source to obtain the identified attribute from the attribute source.

76. The method of claim 75 wherein invoking the attribute source to obtain
1 the identified attribute from the attribute source includes an identification of the attribute.

77. The method of claim 75 wherein invoking the attribute source to obtain
1 the identified attribute from the attribute source includes an indication of a maximum time in
2 which the attribute source is expected to supply the identified attribute.

78. The method of claim 75, further comprising receiving the identified
1 attribute from the attribute source in response to invocation of the attribute source.

79. The method of claim 78 wherein receiving the identified attribute from
1 the attribute source includes a value of the attribute.

80. The method of claim 79 wherein receiving the identified attribute from
1 the attribute source includes an indication of the level of uncertainty of the value of the
2 attribute.

81. The method of claim 79 wherein receiving the identified attribute from
1 the attribute source includes an indication of the time at which the value of the attribute is
2 most accurate.

82. The method of claim 79 wherein receiving the identified attribute from
1 the attribute source includes an indication of the units in which the value of the attribute is
2 expressed.

83. The method of claim 79 wherein receiving the identified attribute from
1 the attribute source includes an indication of the format in which the value of the attribute is
2 expressed.

84. The method of claim 74 further comprising:

receiving an invocation request to provide an attribute value, the invocation request to provide an attribute value identifying the attribute identified by the invocation request to register an attribute, the invocation request to provide an attribute value being generated by a requesting attribute consumer; and

in response to receiving the invocation request to provide an attribute value:

using the generated indication of the specified manner of invoking the attribute source to obtain the identified attribute from the attribute source to invoke the attribute source to obtain the identified attribute from the attribute source, and

providing the attribute value obtained from the attribute source to the requesting attribute consumer.

85. The method of claim 73, wherein the invocation request further specifies a name of the attribute, and wherein the generated indication that the identified attribute can be obtained from the attribute source includes an indication of the specified name of the attribute.

86. The method of claim 73, wherein the invocation request further specifies a name of the attribute source, and wherein the generated indication that the identified attribute can be obtained from the attribute source includes an indication of the specified name of the attribute source.

87. The method of claim 73, wherein the invocation request further specifies a data type of the attribute, and wherein the generated indication that the identified attribute can be obtained from the attribute source includes an indication of the specified data type of the attribute.

88. The method of claim 73, wherein the invocation request further specifies an indication of a format of the attribute, and wherein the generated indication that the identified attribute can be obtained from the attribute source includes an indication of the specified format indication of the attribute.

89. A computer-readable medium whose contents cause a computing device
1 to exchange context attributes by:
3 receiving an invocation request to register an attribute, the invocation request
4 being generated by an attribute source, the invocation request identifying the attribute to be
5 registered; and
6 in response to receiving the invocation request, generating an indication that
7 the identified attribute can be obtained from the attribute source.

90. The computer-readable medium of claim 89, wherein the invocation
1 request further specifies a manner of invoking the attribute source to obtain the identified
2 attribute from the attribute source, and wherein the generated indication that the identified
3 attribute can be obtained from the attribute source includes an indication of the specified
4 manner of invoking the attribute source to obtain the identified attribute from the attribute
5 source.

91. The computer-readable medium of claim 90 wherein the contents of the
1 computer-readable medium further cause the computing device to use the indication of the
2 specified manner of invoking the attribute source to obtain the identified attribute from the
3 attribute source included in the indication that the identified attribute can be obtained from
4 the attribute source to invoke the attribute source to obtain the identified attribute from the
5 attribute source.

92. The computer-readable medium of claim 91 wherein invoking the
1 attribute source to obtain the identified attribute from the attribute source includes an
2 identification of the attribute.

93. The computer-readable medium of claim 90 wherein the contents of the
1 computer-readable medium further cause the computing device to:
3 receive an invocation request to provide an attribute value, the invocation
4 request to provide an attribute value identifying the attribute identified by the invocation

5 request to register an attribute, the invocation request to provide an attribute value being
6 generated by a requesting attribute consumer; and
7 in response to receiving the invocation request to provide an attribute value:
8 use the generated indication of the specified manner of invoking the attribute
9 source to obtain the identified attribute from the attribute source to invoke the attribute
10 source to obtain the identified attribute from the attribute source, and
11 providing the attribute value obtained from the attribute source to the
12 requesting attribute consumer.

94. A computer memory containing an invocation data structure generated
1 by an attribute source, the data structure representing an invocation request to register an
2 attribute, and comprising:
3
4 an indication that the represented invocation request is a request for the
5 invocation of a function to register an attribute;
6 information identifying the attribute to be registered; and
7 information identifying the attribute source,
8 such that the data structure may be used to obtain the registered attribute from
9 the attribute source.

95. The computer memory of claim 94 wherein the data structure further
1 comprises information specifying a manner of invoking the attribute source to obtain the
2 identified attribute from the attribute source,
3
4 such that the data structure may be used to obtain the registered attribute from
5 the attribute source by invoking the attribute source in the specified manner.

96. The computer memory of claim 94, wherein the data structure further
1 comprises information specifying a name of the attribute.

97. The computer memory of claim 94, wherein the data structure further
1 comprises information specifying a name of the attribute source.

98. A method in a computing device for exchanging context attributes,

1 comprising:

3 receiving an invocation request to create a condition based on an attribute, the
4 invocation request identifying the attribute on which the condition is based, the invocation
5 request specifying a test to perform on the identified attribute; and

6 in response to receiving the invocation request, generating an indication of the
7 condition, the generated indication identifying the attribute identified by the invocation
8 request and indicating the specified test.

99. The method of claim 98 wherein the invocation request further specifies

1 a name of the condition, and wherein the generated indication of the condition further
2 indicates the specified condition name.

100. The method of claim 99 further comprising:

2 receiving an invocation request to evaluate a condition, the invocation request
3 to evaluate a condition being generated by a requesting condition consumer, the invocation
4 request to evaluate a condition specifying the condition name specified by the invocation
5 request to create a condition; and

6 in response to receiving the invocation request to evaluate a condition:

7 using the generated indication of the condition to evaluate the condition based
8 upon the attribute, and

9 providing to the requesting condition consumer an indication of the value of
10 the condition as evaluated.

101. A method in a computing device for exchanging context attributes,

1 comprising:

3 receiving an invocation request to create a condition monitor based on a
4 condition, the invocation request identifying the condition on which the condition monitor is
5 based, the invocation request specifying a test to perform on the identified condition monitor,
6 the invocation request being generated by a requesting condition monitor consumer; and

7 in response to receiving the invocation request, generating an indication of the
8 condition monitor, the generated indication identifying the condition identified by the
9 invocation request and indicating the specified test.

102. The method of claim 100 further comprising:

2 at a time after receiving the invocation request, determining that the specified
3 test has been satisfied; and

4 in response to determining that the specified test has been satisfied, notifying
5 the requesting condition monitor consumer that the specified test has been satisfied.

103. The method of claim 100 wherein the invocation request further
1 specifies a frequency with which the specified test is to be evaluated, and wherein the
2 method further comprises reiteratively evaluating the specified test at substantially the
3 specified frequency to determine whether the specified test has been satisfied

104. The method of claim 100 wherein the identified condition may
1 transition between values of true and false, and wherein the specified test is either (1) that the
2 identified condition transitions from the value false to the value true, (2) that the identified
3 condition transitions from the value true to the value false, or (3) that the identified condition
4 transitions either from the value false to the value true or from the value true to the value
5 false.

105. The method of claim 100 wherein the invocation request further
1 specifies a manner of invoking the requesting condition monitor consumer to notify the
2 requesting condition monitor consumer that the specified test has been satisfied, and wherein
3 the generated indication of the condition monitor includes an indication of the specified
4 manner of invoking the requesting condition monitor consumer to notify the requesting
5 condition monitor consumer that the specified test has been satisfied.

106. The method of claim 100 wherein the invocation request further
1 specifies a name of the condition monitor, and wherein the generated indication of the
2 condition further indicates the specified condition monitor name.

107. The method of claim 106 further comprising:

after receiving an invocation request to create a condition monitor, receiving an invocation request to suspend the operation of a condition monitor, the invocation request to suspend the operation of the condition monitor specifying the name of the condition monitor specified by the invocation request to create a condition monitor; and

in response to receiving the invocation request to suspend the operation of a condition monitor, modifying the indication of the condition monitor to indicate that operation of the condition monitor is suspended.

108. The method of claim 107, further comprising:

after receiving an invocation request to suspend the operation of a condition monitor, receiving an invocation request to resume the operation of a condition monitor, the invocation request to resume the operation of a condition monitor specifying the name of the condition monitor specified by the invocation request to create a condition monitor; and

in response to receiving the invocation request to suspend the operation of a condition monitor, modifying the indication of the condition monitor to indicate that operation of the condition monitor is not suspended.

109. The method of claim 108, further comprising:

during the period after receiving an invocation request to create a condition monitor and before receiving an invocation request to suspend the operation of a condition monitor, periodically evaluating the specified test;

during the period after receiving an invocation request to suspend the operation of a condition monitor and before receiving an invocation request to resume the operation of a condition monitor, omitting to evaluate the specified test; and

after receiving an invocation request to resume the operation of a condition monitor, periodically evaluating the specified test.

110. A method in a computing device for exchanging context attributes,

comprising:

3 receiving an invocation request to provide an indication of attribute consumers
4 that are dependent on an attribute, the invocation request identifying the attribute about
5 which an indication is to be provided, the identified attribute having a source, the invocation
6 request being generated by the source of the identified attribute;

7 in response to receiving the invocation request, providing an indication of
8 attribute consumers that are dependent on the identified attribute.

111. The method of claim 110 wherein the provided indication of attribute
1 consumers that are dependent on the identified attribute is the number of attribute consumers
2 that are dependent on the identified attribute.

112. The method of claim 110 wherein the provided indication of attribute
1 consumers that are dependent on the identified attribute is an indication of the identities of
2 attribute consumers that are dependent on the identified attribute.

113. A method in a computing device for exchanging context attributes,
1 comprising:

3 receiving an invocation request to launch a context server from which values of
4 one or more attributes can be obtained, the invocation request identifying the context server
5 to be launched, the invocation request being generated by a context client; and

6 in response to receiving the invocation request, launching the identified context
7 server.

114. A method in a computing device for exchanging context attributes,
1 comprising:

3 receiving an invocation request to provide an attribute value, the invocation
4 request being generated by a requesting attribute consumer, the invocation request
5 identifying the attribute whose value is to be provided; and

6 in response to receiving the invocation request, providing a value for the
7 identified attribute to the requesting attribute consumer.

115. One or more computer memories collectively containing an attribute
1 instance registration data structure reflecting a set of registered attribute instances, each
2 attribute instance having a source having an identity, the data structure comprising a
3 multiplicity of entries, each entry corresponding to a registered available attribute instance
4 and containing information indicating:

6 an attribute name for the registered attribute instance; and

7 the identity of the source of the registered attribute instance,

8 such that the data structure may be used to determine the attribute names and sources of
9 registered attribute instances.

116. The computer memories of claim 115 wherein a plurality of the
1 multiplicity of entries each contain an indication of an earlier-retrieved value of the attribute
2 instances to which they correspond.

117. The computer memories of claim 116 wherein the plurality of entries
1 each further contain an indication of a time at which the earlier-retrieved value of the
2 attribute instance was effective.

118. The computer memories of claim 115 wherein a plurality of the
1 multiplicity of entries each contain an indication of a manner of invoking the source of the
2 corresponding attribute instance to obtain a value for the attribute instance from the source

119. One or more computer memories collectively containing a context
1 attribute condition data structure, the data structure comprising a plurality of entries, each
2 entry corresponding to a condition based upon an attribute and containing:

4 information identifying the attribute upon which the condition is based; and

5 information specifying the test to be performed on the identified attribute in
6 order to evaluation the condition,

7 such that the data structure may be used to evaluate the conditions to which its entries
8 correspond.

120. The computer memories of claim 119 wherein each entry further
1 contains an indication of a name of the corresponding condition,
3 such that the contents of the data structure may be used to evaluate a condition having a
4 particular name.

121. One or more computer memories collectively containing a context
1 attribute condition monitor data structure, the data structure comprising a plurality of entries,
2 each entry corresponding to a condition monitor based upon condition and containing:
4 information identifying the condition upon which the condition monitor is
5 based; and
6 information specifying the test to be performed on the identified condition in
7 order to determine whether the condition monitor has been triggered,
8 such that the data structure may be used to determine whether the conditions monitors to
9 which its entries correspond have been triggered.

122. The computer memories of claim 121 wherein each entry further
1 specifies a frequency at which the corresponding condition monitor is to be evaluated to
2 determine whether the condition monitor has been triggered,
4 such that the contents of the data structure may be used to evaluate the condition monitors to
5 which its rows correspond each at the frequency specified for it.

123. The computer memories of claim 121 wherein each entry further
1 contains an indication of whether operation of the corresponding condition monitor is
2 suspended,
4 such that the contents of the data structure may be used to determine which of the condition
5 monitors to which its rows correspond need not be evaluated.

124. The method of claim 1, further comprising the step of, in response to
2 receiving the invocation request, providing a name and a value of a property associated with
3 the identified attribute to the requesting attribute consumer.

1 125. The method of claim 6, furthering comprising obtaining a value of a
2 property associated with the identified attribute from the attribute source with which the
3 identified attribute is associated.

1 126. The method of claim 125, furthering comprising providing the obtained
2 associated property value to the requesting attribute consumer.

1 127. The method of claim 125, wherein the obtained property is security
2 information associated with the identified attribute, further comprising using the security
3 information to determine whether to provide a value for the identified attribute to the
4 requesting attribute consumer.

1 128. The computer memories of claim 116 wherein a selected entry contains
2 information indicating a value of a property associated with the attribute instance to which
3 the entry corresponds.

1 129. The computer memories of claim 128 wherein the selected entry
2 contains information indicating a name of the property associated with the attribute instance
3 to which the selected entry corresponds.

1 130. The computer memories of claim 128 wherein the property value
2 constitutes security information associated with the attribute instance to which the selected
3 entry corresponds,
4 such that the data structure may be used to determine whether to provide a value of the
5 attribute instance to which the selected row corresponds to a particular attribute consumer.

6

INTERFACE FOR EXCHANGING CONTEXT DATA

ABSTRACT

A facility for exchanging context attributes is described. A characterization module receives an invocation request to provide an attribute value that was generated by a requesting attribute consumer. The received invocation request identifies the attribute whose value is to be provided. In response to receiving the invocation request, the characterization module provides a value for the identified attribute to the requesting attribute consumer.

8005-fig1 vsd




Fig. 1

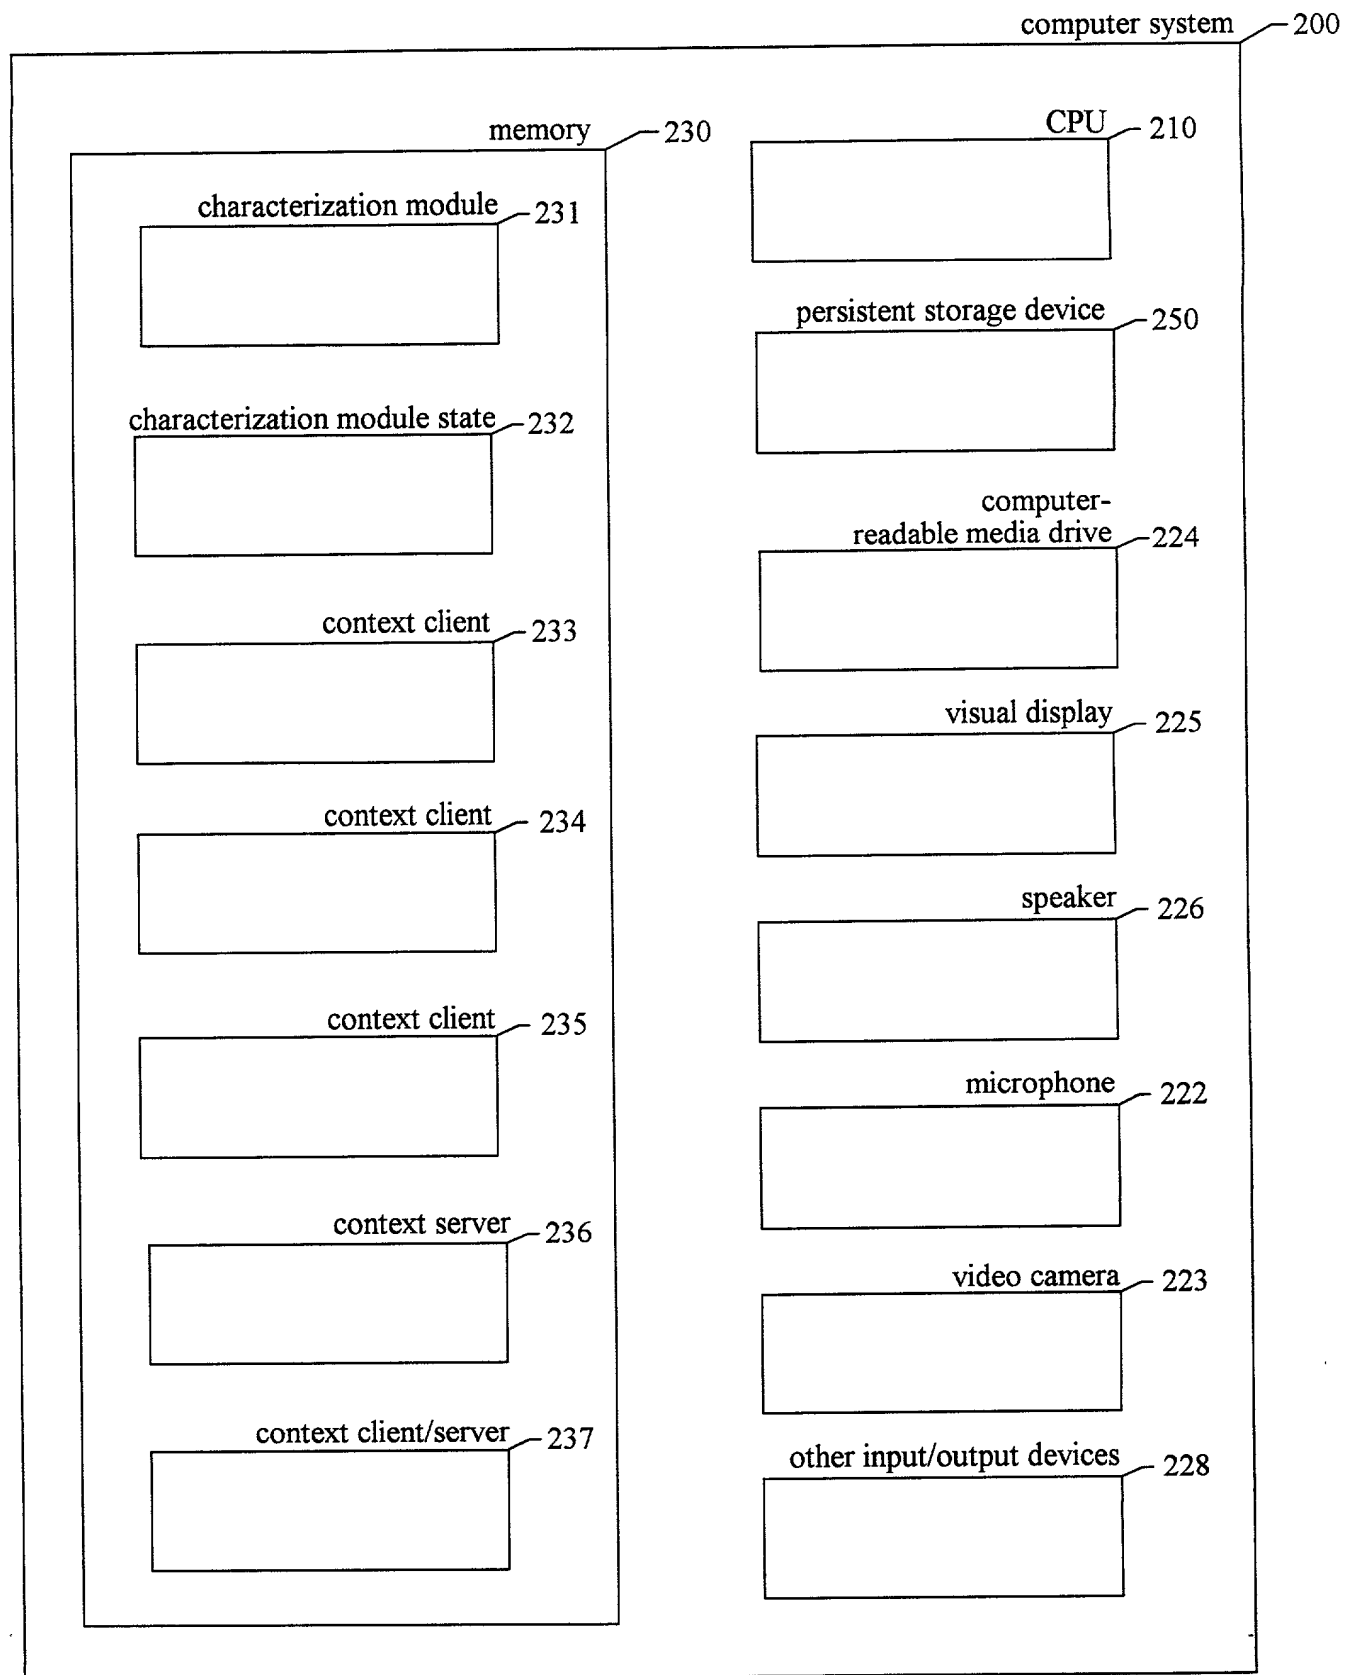


Fig. 2




Fig. 3

context server table					400
context server name	version	installation date	filename	request handler	415
401—gps	1.00.315	2/10/2000	gps.exe	(reference)	
402—ips	1.00.315	2/21/2000	ips.exe	(reference)	
403—location_region_analysis	1.00.315	2/10/2000	l_r_a.exe	(reference)	
411		412	413	414	415

Fig. 4

attribute instance table						500
attribute name	context server name	value	uncertainty	timestamp	units	516
501—user.location	gps	47°38.73' N, 122° 18.43' W	0° .09'	13:11:04.023 2/22/2000	degrees/minutes	
502—user.location	ips	47°38.745' N, 122° 18.424' W	0° .021'	13:11:01.118 2/22/2000	degrees/minutes	
503—user.elevation	ips	22	.5	13:11:01.118 2/22/2000	meters	
504—user.in_region	location_region_analysis	none	none	none	none	515
511		512	513	514	515	516

Fig. 5

context client table		600
context client name	message handler	
601 location_map		
602 distance_from_home	(reference)	
603 region_analysis	(reference)	
	(reference)	

Fig. 6

attribute or instance registration table			700
context client name	attribute name	context server name	
701 region_analysis	user.location		
702 distance_from_home	user.location		
703 location_map	user.elevation	ips	
704 location_map	user.in_region	region_analysis	

Fig. 7

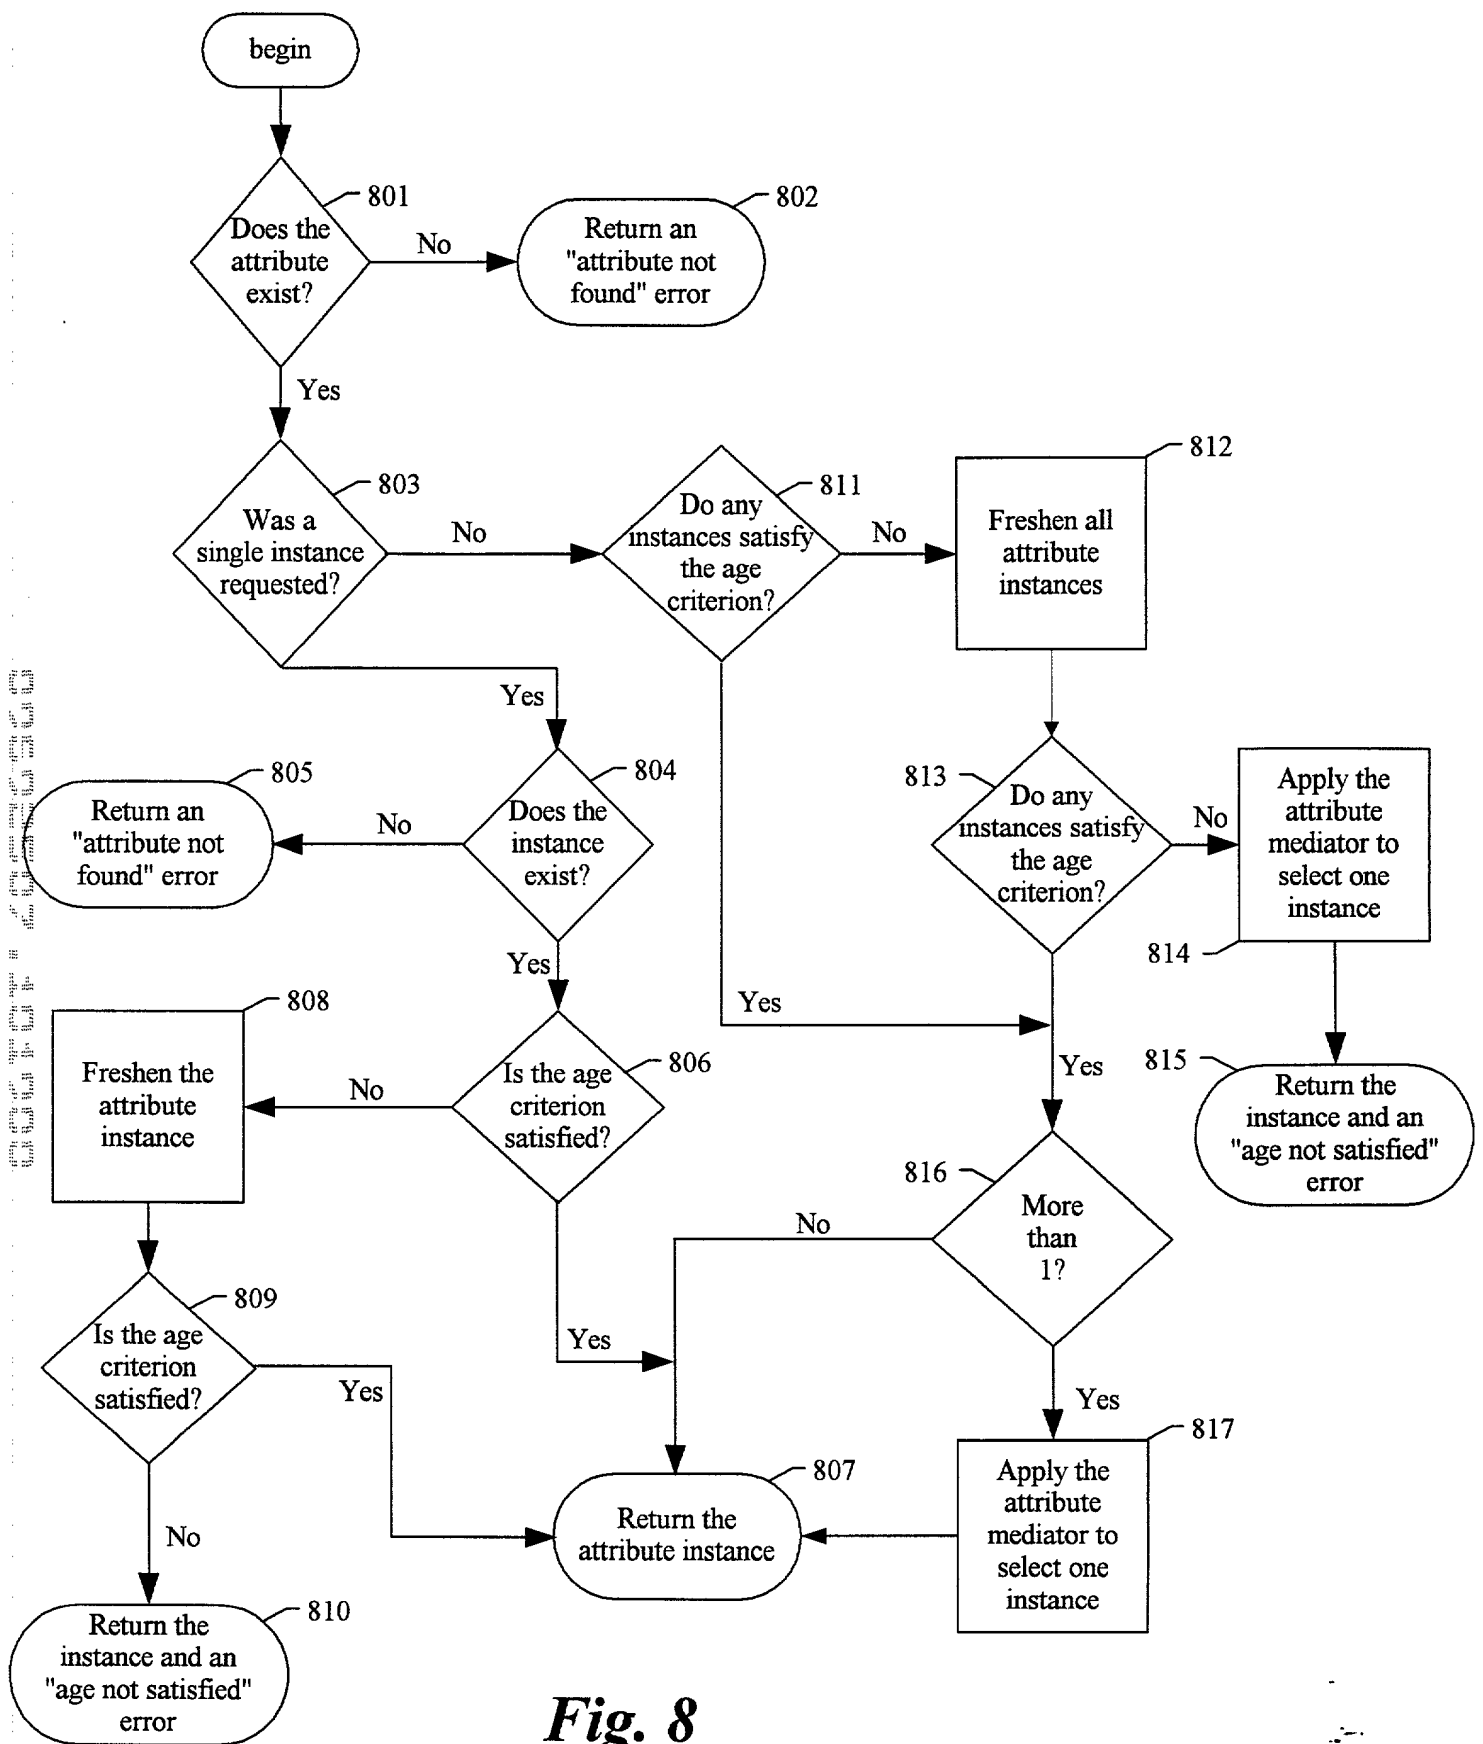


Fig. 8

attribute name	context server name	value	uncertainty	timestamp	units	number of context clients consuming
user.location	gps	47°38.73' N, 122° 18.43' W	0° .09'	13:11:04.023 2/22/2000	degrees/minutes	2
user.location	ips	47°38.745' N, 122° 18.424' W	0° .021'	13:11:01.118 2/22/2000	degrees/minutes	2
user.elevation	ips	22.25	.5	13:11:06.565 2/22/2000	meters	1
user.in_region	location_analysis	none	none	none	none	1

Fig. 9

condition name	context client name	first logical parameter	second logical parameter	comparison value	logical operator
in_region_true	region_analysis	user.in_region	none	TRUE	=

Fig. 10

condition monitor name	context client name	condition name	behavior	frequency	condition last evaluated	trigger handler reference	stopped
region_boundary_crossed	region_analysis	in_region_true	any state change	30	13:11:29.101 2/22/2000	(reference)	no

Fig. 11

condition monitor name	context client name	condition name	behavior	frequency	condition last evaluated	trigger handler reference	stopped
region_boundary_crossed	region_analysis	in_region_true	any state change	30	13:11:59.101 2/22/2000	(reference)	yes

Fig. 12

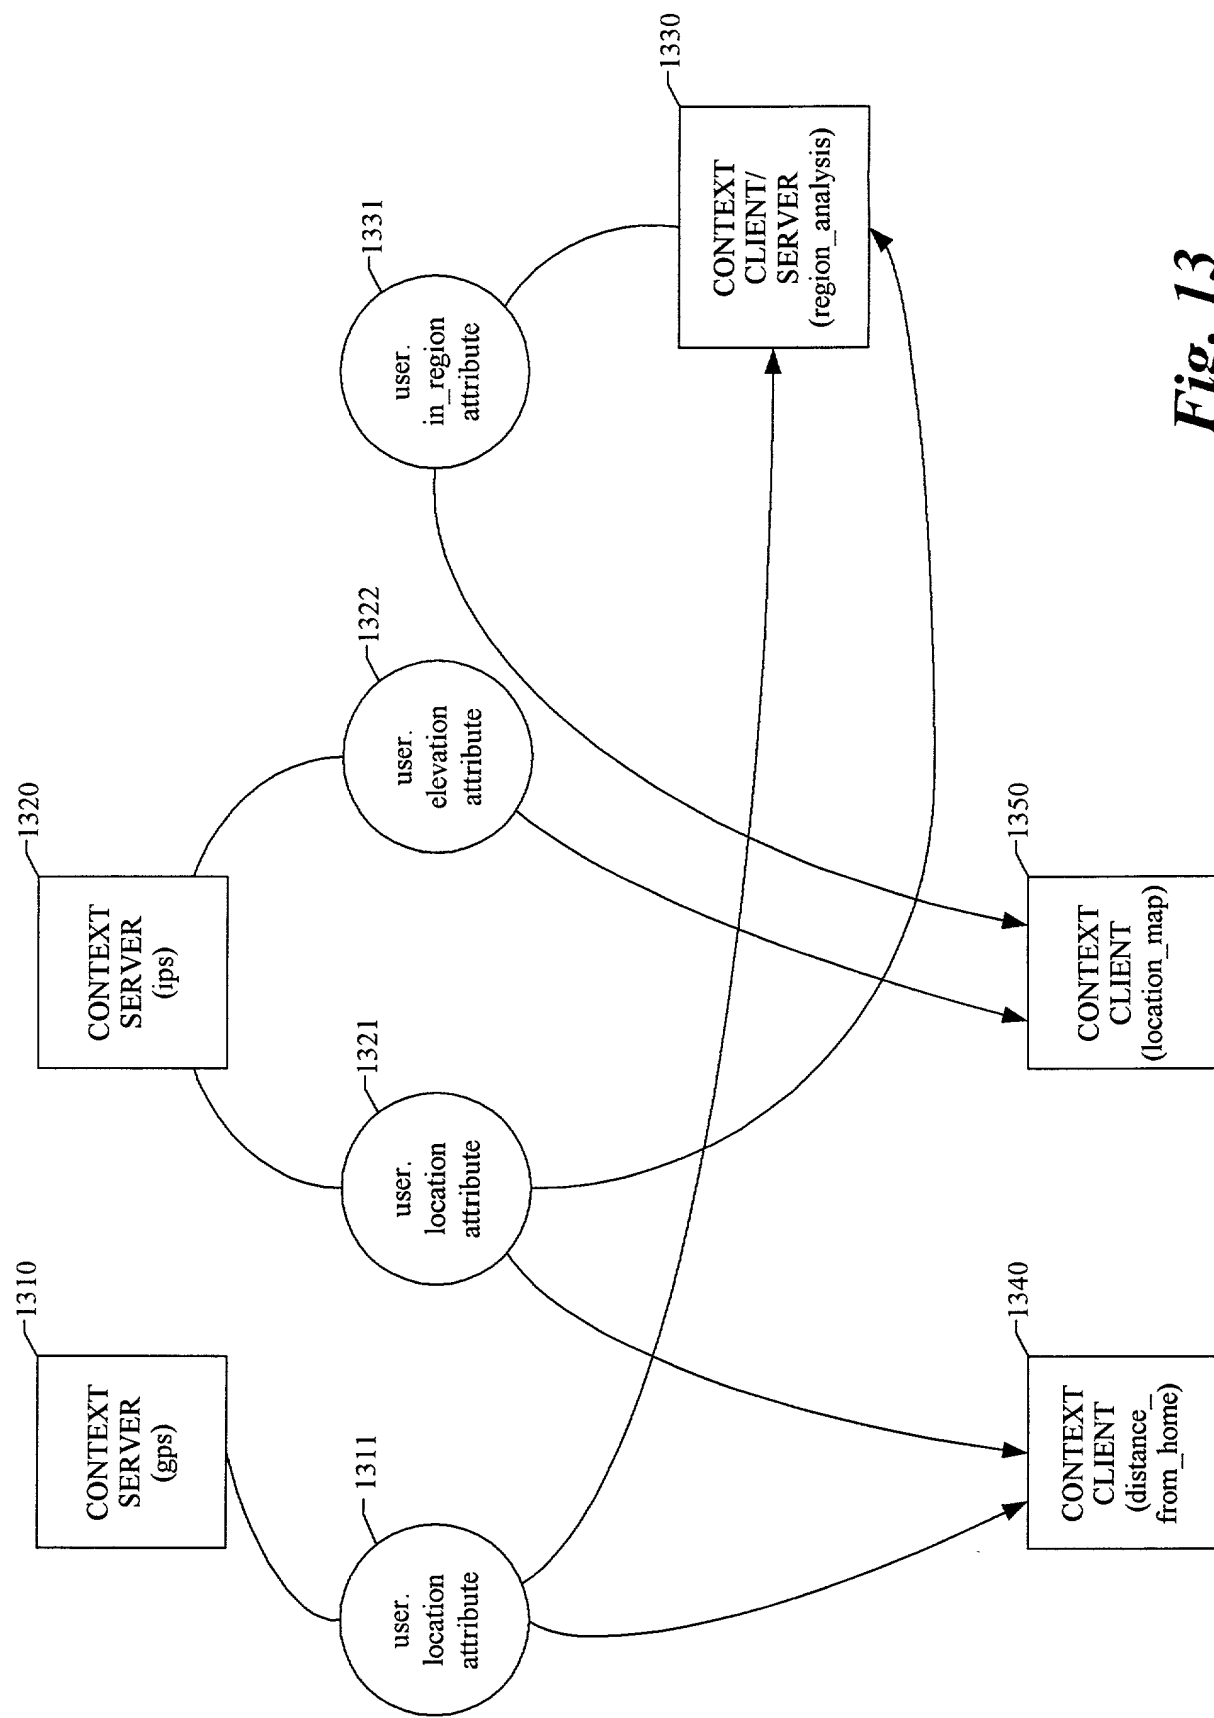


Fig. 13

any other data may be included in the data structure for the user location

attribute instance property table				1400
attribute name	context server name	property name	property value	
user location	ips	security_token	5A637AR	

1401

1401

1402

1403

1404

Fig. 14